
The Three P's of Total Risk Management

Andrew W. Lo

Current risk-management practices are based on probabilities of extreme dollar losses (e.g., measures like Value at Risk), but these measures capture only part of the story. Any complete risk-management system must address two other important factors—prices and preferences. Together with probabilities, these compose the three P's of "Total Risk Management." This article describes how the three P's interact to determine sensible risk profiles for corporations and for individuals—guidelines for how much risk to bear and how much to hedge. By synthesizing existing research in economics, psychology, and decision sciences and through an ambitious research agenda to extend this synthesis into other disciplines, a complete and systematic approach to rational decision making in an uncertain world is within reach.

Although rational decision making in the face of uncertainty is by no means a new aspect of the human condition,¹ recent events have helped to renew and deepen interest in risk management. Two forces in particular have shaped this trend: advances in financial technology (models for pricing derivative instruments and computationally efficient means for implementing them) and an ever-increasing demand for new and exotic financial engineering products (perhaps because of increased market volatility or simply because of the growing complexity of the global financial system). These forces, coupled with such recent calamities as those of Orange County, Gibson Greetings, Metallgesellschaft, Procter & Gamble, and Barings Securities, provide more than sufficient motivation for a thriving risk-management industry.

Current risk-management practices focus almost exclusively on the *statistical* aspects of risk. For example, one of the most popular risk-management tools, Value at Risk (VAR), is described in J.P. Morgan's RiskMetrics system documentation in the following way:

Value at Risk is an estimate, with a predefined confidence interval, of how much one can lose from holding a position over a set horizon. Potential horizons may be one day for typical trading activities or a month or longer for portfolio management. The methods described in our documentation use historical returns to forecast volatilities and correlations that are

then used to estimate the market risk. These statistics can be applied across a set of asset classes covering products used by financial institutions, corporations, and institutional investors. (Morgan Guaranty Trust Company, 1995, p. 2)

Although measures like VAR play an important role in quantifying risk exposure, they address only one piece of the risk-management puzzle—probabilities. Probabilities are an indispensable input into the risk-management process, but they do not determine how much risk a corporation should bear and how much should be hedged. In this article, I argue that any complete risk-management protocol—what might be called "Total Risk Management" (TRM),² to borrow a phrase from the quality control literature—must include two other pieces: prices and preferences. Together with probabilities, these three P's form the basis of a systematic approach to rational decision making in an uncertain world. All three P's are central to TRM: prices in considering how much one must pay for hedging various risks; probabilities for assessing the likelihood of those risks; and preferences for deciding how much risk to bear and how much to hedge.

Despite being a trendy catchphrase, TRM has deep intellectual roots in economics, statistics, and mathematics and is based on research that can be traced back to the very foundations of probability theory (Ramsey 1926), statistical inference (Savage 1954), and game theory (von Neumann and Morgenstern 1944). Of course, the term "risk management" never appears in that literature, but the issues that these early pioneers grappled with are precisely those that concern us today. Indeed, I

Andrew W. Lo is Harris & Harris Group Professor at the Massachusetts Institute of Technology's Sloan School of Management.

hope to show much can be gained by synthesizing and extending the various disparate strands of research that have grown out of these seminal works: Current risk-management practices have so far drawn on only one such strand.

The Three P's

To understand the interactions among prices, probabilities, and preferences, consider the most fundamental principle of economics, namely, the law of supply and demand. This law states that the market price of any commodity and the quantity traded are determined by the intersection of supply and demand curves, where the demand curve represents the schedule of quantities desired by consumers at various prices and the supply curve represents the schedule of quantities producers are willing to supply at various prices. The intersection of these two curves is the price–quantity pair that satisfies both consumers and producers; any other price–quantity combination may serve one group's interests but not the other's.

Even in such an elementary description of a market, the three P's are present. The demand curve is the aggregation of individual consumers' demands, each derived from optimizing an individual's preferences, subject to a budget constraint that depends on prices and other factors (e.g., income, savings requirements, and borrowing costs). Similarly, the supply curve is the aggregation of individual producers' outputs, each derived from optimizing an entrepreneur's production function, subject to a resource constraint that also depends on prices and other factors (e.g., costs of materials, wages, and trade credit). And probabilities affect both consumers and producers as they formulate their consumption and production plans over time and in the face of uncertainty—uncertain income, uncertain costs, and uncertain business conditions.

Formal models of asset prices and financial markets, such as those of Merton (1973b), Lucas (1978), Breeden (1979), and Cox, Ingersoll, and Ross (1985), show precisely how the three P's simultaneously determine an "equilibrium" in which demand equals supply across all markets in an uncertain world where individuals and corporations act rationally to optimize their own welfare. Typically, these models imply that a security's price is equal to the present value of all future cash flows to which the security's owner is entitled. Two aspects make this calculation unusually challenging: Future cash flows are uncertain, and so are discount rates. Although pricing equations that account for both aspects are often daunting,³ their

intuition is straightforward and follows from the dividend discount formula: Today's price must equal the expected sum of all future dividends multiplied by discount factors that act as "exchange rates" between dollars today and dollars at future dates. If prices do not satisfy this condition, then there must be a misallocation of resources between today and some future date. This situation would be tantamount to two commodities selling for different prices in two countries after exchange rates have been taken into account.

What determines the exchange rate? For individuals, it is influenced by their preferences (the ratio of marginal utilities of consumption, to be precise), and it is determined in an equilibrium by the aggregation of all the preferences of individuals in the market through the equalization of supply and demand.

These models show that equilibrium is a powerful concept that provides a kind of adding-up constraint for the three P's: In an equilibrium, any two P's automatically determine the third. For example, given an equilibrium in which preferences and probabilities are specified, prices are determined exactly (this is the central focus of the asset-pricing literature in economics). Alternatively, given an equilibrium in which prices and probabilities are specified, preferences can be inferred exactly (see, for example, Bick 1990, He and Leland 1993, Ait-Sahalia and Lo 1998b, and Jackwerth 1998). And given prices and preferences, probabilities can be extracted (see, for example, Rubinstein 1994 and Jackwerth and Rubinstein 1996).

This functional relationship suggests that the three P's are inextricably linked, and even though current risk-management practices tend to focus on only one or two of them, all three P's are always present and their interactions must be considered carefully. In the sections to follow, I consider each of the three P's in turn and describe how each is related to the other two. Although all three P's are crucial for any TRM system, I will argue that preferences may be the most fundamental, the least understood, and therefore, the most pressing challenge for current risk-management research.

Prices

One of the great successes of modern economics is the subfield known as asset pricing,⁴ and within asset pricing, surely the crowning achievement in the past half-century is the development of precise mathematical models for pricing and hedging derivative securities. The speed with which the ideas of Black and Scholes (1973) and Merton (1973a) were embraced, both in academia and in industry, is

unprecedented among the social sciences, and this, no doubt, has contributed to the broad success of risk-management policies and technologies.

The asset-pricing literature is so deep and rich that there is little need to expound on the importance of prices for risk management. Nevertheless, even for this most studied of the three P's, some subtle links to the other two P's are worth explicating.

Perhaps the most important insight of the Black–Scholes–Merton framework is that under certain conditions, the frequent trading of a small number of long-lived securities can create new investment opportunities that would otherwise be unavailable to investors. These conditions—now known collectively as “dynamic spanning” or “dynamic market completeness”—and the asset-pricing models on which they are based have generated a rich literature, and an even richer industry, in which complex financial securities are synthetically replicated by sophisticated trading strategies involving considerably simpler instruments.⁵ This approach lies at the heart of the celebrated Black–Scholes–Merton option-pricing formula and, more generally, the no-arbitrage method of pricing and hedging other derivative securities.

The success of derivative-pricing models is central to risk management for at least two reasons. The first reason is obvious: Complex derivative securities, on which most risk-management practices are built, can be priced accurately and hedged effectively using the Black–Scholes–Merton methodology and its extensions.

The second reason is considerably more subtle and can be best understood through a paradox. The accuracy of derivative-pricing models seems to be at odds with the framework discussed in the beginning of this section in which the three P's were said to be inseparable. In particular, in typical derivative-pricing models (those based on continuous-time stochastic processes and the usual partial differential equations), prices and probabilities are featured prominently but no mention is made of investors' preferences. Indeed, such models are often trumpeted as being “preference free,” being based solely on arbitrage arguments and not on equilibrium or supply-and-demand considerations. In fact, the risk preferences of individual investors never enter into the Black–Scholes formula: As long as the Black–Scholes assumptions hold (and these assumptions do not seem to restrict preferences in any way), a retired widow living on social security places the same value on a call option as a 25-year-old, unmarried bond trader! If derivatives are priced solely by arbitrage, where is the third P in derivative-pricing models?

The answer to this paradox lies in the fact that

preferences do enter the Black–Scholes formula but in a subtle and indirect way. In particular, the assumption that the underlying asset's price dynamics are governed by a particular stochastic process—typically, geometric Brownian motion—restricts the type of possible preferences (see, e.g., Bick 1990 and He and Leland 1993).

Moreover, the parameters of the stochastic process (e.g., the drift and diffusion coefficients in geometric Brownian motion) are determined in equilibrium, not by arbitrage. After all, the drift of the underlying asset's price process is the asset's instantaneous expected return, and one of the basic tenets of modern finance is that expected returns and risk are jointly determined by supply and demand (see, in particular, Sharpe 1964 and Merton 1973b). This intuition applies even though the drift does not appear in derivative-pricing formulas, because the drift and diffusion coefficients are linked (see, e.g., Lo and Wang 1995), and it is telling that the original Black and Scholes derivation used equilibrium arguments to arrive at their celebrated partial differential equation.⁶ In more-complex derivative-pricing models, such as those in which perfect replication is not possible—the case of stochastic volatility, for example—equilibrium arguments must be used explicitly to derive the pricing equation.

Therefore, although derivative-pricing formulas may seem preference free, they do contain implicit assumptions about preferences and probabilities. The three P's are inextricably linked even in arbitrage-based pricing models.

Probabilities

Through the centuries, researchers have proposed a number of approaches to modeling and decision making in an uncertain world—astrology, numerology, and reading animal entrails, to name just a few—but none has enjoyed as much success as the mathematical theory of probability. The concept of randomness can be traced back to the Greeks, but formal and numerical notions of probability did not arise until the 17th century in the context of games of chance.⁷ Since then, probability theory has developed into a rich and deep discipline that has become central to virtually every scientific discipline, including financial economics and risk management.

As with prices, probabilities are fairly well understood. We are familiar with the algebra of probabilities—the fact that probabilities are non-negative and sum to 1, that the probability of two independent events occurring simultaneously is the product of the two events' probabilities, and so on. We understand the mathematics of probability

distributions, the critical role that correlation plays in risk management, and the sensitivity of VAR and other risk-management tools to “tail” probabilities (the probabilities associated with rare but potentially ruinous events).

But one important aspect of probabilities has been largely ignored in the risk-management literature: the distinction between “objective” and “subjective” probabilities, usually attributed to the 18th century mathematician James Bernoulli. Objective probabilities, also called “statistical” or “aleatory” probabilities, are based on the notion of relative frequencies in repeated experiments (e.g., coin tosses, rolls of the dice). Such probabilities have clear empirical origins—the probability of rolling a six is $1/6$, and this fact can be verified by rolling a fair die many times and computing the ratio of sixes to the total number of trials. The probability $1/6$ depends on the nature of the experiment, not on the characteristics of the experimenter—hence, the term “objective” probabilities.

On the other hand, subjective probabilities, also called “personal” or “epistemic” probabilities, measure “degrees of belief,” which need not be based on statistical phenomena, such as repeated coin tosses. For example, the event “There is intelligent life on other planets” cannot be given a relative frequency interpretation—we cannot conduct repeated trials of this event. Nevertheless, we can easily imagine an individual possessing a certain level of conviction about the likelihood of such an event. This level of conviction can be interpreted as a kind of probability, a *subjective* one that can differ from one individual to another. Subjective probability is a powerful concept that extends the reach of probability theory to a much broader set of applications, many of which are central to risk management. In particular, one of the most critical aspects of any risk-management protocol is the ability to assess the likelihood of and prepare for events that may never have occurred in the past (e.g., the unprecedented global flight to quality by financial market participants during August 1998 and the surprising degree of correlation between yield spreads, exchange rates, and commodity and stock prices that it created).

The link between subjective probabilities and risk management becomes even stronger when considered in light of the foundations on which subjective probabilities are built. The three main architects of this theory—Ramsey, De Finetti (1937), and Savage—argued that, despite the individualistic nature of subjective probabilities, they must still satisfy the same mathematical laws as objective probabilities; otherwise, arbitrage oppor-

tunities will arise.⁸ For example, consider the basic axiom of objective probability: The probability of any event H and the probability of its complement “not H ,” denoted by H^C , must sum to one—that is,

$$\text{prob}(H) + \text{prob}(H^C) = 1 \quad (1)$$

—because H and H^C are mutually exclusive and exhaustive; in other words, only one or the other will occur, and together, these two events cover all possible outcomes. Equation 1 can be readily verified for objective probabilities by applying simple arithmetic to relative frequencies, but can it be “proved” for subjective probabilities as well? In other words, must individuals’ degrees of belief also satisfy this basic property of objective probabilities? The answer—conjectured by Ramsey and proven vigorously by De Finetti and Savage—is yes, if arbitrage opportunities or “free lunches” are ruled out.

To see why, consider an individual who attaches a probability of 50 percent to an event H and 75 percent to its complement H^C , clearly a violation of Equation 1. Such subjective probabilities imply that such an individual would be willing to take a bet at even odds that H occurs and, at the same time, would also be willing to take a bet at 3:1 odds that H^C occurs. Someone taking the other side of these two bets—placing \$50 on the first bet and \$25 on the second—would have a total stake of \$75 but be assured of receiving \$100 regardless of the outcome, yielding a riskless profit of \$25—an arbitrage! De Finetti proved that the only set of odds for which such an arbitrage cannot be constructed is one in which Equation 1 and the other basic axioms of probability theory are satisfied. Therefore, despite the fact that subjective probabilities measure only degrees of belief and are not based on relative frequencies, they behave like objective probabilities in every respect. This principle is often called the “Dutch book theorem,” an allusion to a kind of arbitrage transaction known as a “Dutch book.”

The relationship between subjective probabilities and risk management is clear: Probability assessments, particularly those of rare events or events that have never occurred, must be internally consistent; otherwise, prices derived from such probabilities may be inconsistent, which leads to arbitrage opportunities for others. More importantly, decisions based on inconsistent probabilities can lead to significant financial losses and unintended risk exposures.

The Dutch book theorem also shows that prices and probabilities are related in a profound way and that neither can be fully understood in isolation and without reference to the other. But this leaves open

the question of how subjective probabilities are determined. The answer—proposed by Savage—is the third and most important of the three P's of risk management: preferences.

Preferences

Models of individual preferences have their historical roots in the school of social philosophy known as Utilitarianism, a system of ethics proposed in the late 18th century by Jeremy Bentham and James Mill (father of political economist John Stuart Mill) in which the goal of all actions is to maximize general utility or happiness. Although moral philosophers and political theorists have debated the merits of Utilitarianism for more than two centuries, economists were quick to adopt the principle that individuals maximize their utility subject to a budget constraint, with utility defined as any quantitative index of happiness satisfying certain basic properties.

The importance of utility to classical economists sprang from their attempt to define the value of a commodity and to distinguish value from the commodity's market price. In making this distinction, Adam Smith (1776) proposed his now-famous comparison of water and diamonds:

The word *value*, it is to be observed, has two different meanings, and sometimes expresses the utility of some particular object, and sometimes the power of purchasing other goods which the possession of that object conveys. The one may be called "value in use"; the other, "value in exchange." The things which have the greatest value in use have frequently little or no value in exchange; and, on the contrary, those which have the greatest value in exchange have frequently little or no value in use. Nothing is more useful than water; but it will purchase scarce any thing; scarce any thing can be had in exchange for it. A diamond, on the contrary, has scarce any value in use; but a very great quantity of other goods may frequently be had in exchange for it. (p. 147)

By distinguishing "value in exchange" (price) from "value in use" (utility), Smith laid the foundation for the law of supply and demand and the notion of market equilibrium, perhaps the single most important contribution of classical economics. Moreover, Samuelson (1947), in *Foundations of Economic Analysis*, which is largely responsible for much of what is now standard microeconomics, wrote:⁹

It so happens that in a wide number of economic problems it is admissible and even mandatory to regard our equilibrium equations as maximizing (minimizing) conditions. A large part of entrepreneurial behavior is directed towards

maximization of profits with certain implications for minimization of expenditure, etc. Moreover, it is possible to derive operationally meaningful restrictive hypotheses on consumers' demand functions from the assumption that consumers behave so as to maximize an ordinal preference scale of quantities of consumption goods and services. (Of course, this does not imply that they behave rationally in any normative sense.) (Chapter III, pp. 21–22)

The notion of utility can also be extended to cover uncertain outcomes, and the first attempt to do so—in 1738 by Daniel Bernoulli—predated Utilitarianism by almost half a century. Bernoulli's motivation had little to do with utility per se; he was attempting to resolve the St. Petersburg paradox. In this paradox, an individual is offered the following gamble: A fair coin is tossed until it comes up heads, at which point the individual is paid a prize of $\$2^k$, where k is the number of times the coin is tossed. How much should an individual pay for such a gamble? Because the probability of tossing heads for the first time on the k th flip is $1/2^k$, the expected value of this gamble is infinite; yet individuals are typically only willing to pay between \$2 and \$4 to play, which is the paradox. Bernoulli (1738) resolved this paradox by asserting that gamblers do not focus on the expected gain of a wager but, rather, on the expected logarithm of the gain, in which case, the "value in use" of the St. Petersburg gamble is

$$\sum_{k=1}^{\infty} 2^{-k} \log(2^k) = 2 \log 2 \approx 4, \quad (2)$$

a value more consonant with casual empirical observation than the expected value of the gamble.

Although Bernoulli did not present his resolution of the St. Petersburg paradox in terms of utility, the essence of his proposal is to replace expected value as the gambler's objective with expected utility, where utility is defined to be the logarithm of the gain. This approach to decision making under uncertainty was remarkably prescient; it anticipated von Neumann and Morganstern's and Savage's axiomatic derivation of expected utility by more than two centuries. In the framework proposed by these later writers, any individual's preferences can be represented numerically by a utility function $U(X)$ if those preferences satisfy certain axioms.¹⁰ In other words, if an individual's preferences satisfy these axioms—the axioms of expected utility—then a utility function $U(X)$ can be constructed in such a way that the individual's choices among various alternatives will coincide with those choices that maximize the individual's

expected utility, $E[U(X)]$.

Formally, given any two gambles with random payoffs X_1 and X_2 , an individual satisfying the axioms of expected utility will prefer X_1 to X_2 if and only if $E[U(X_1)]$ is greater than $E[U(X_2)]$ for some function $U(\cdot)$ that is unique to each individual. Under these axioms, the function $U(\cdot)$ is a *complete* representation of an individual's preferences—all his or her decisions can be fully delegated to another party under the simple dictum "maximize my expected utility $E[U(X)]$." This powerful representation lies at the heart of virtually every modern approach to pricing financial assets, including modern portfolio theory, mean-variance optimization, the capital asset pricing model, the intertemporal capital asset pricing model, and the Cox–Ingersoll–Ross term-structure model. Expected utility is also central to risk management because the final outcome of any risk-management protocol is a decision about how much risk to bear and how much to hedge; although prices and probabilities surely influence this decision, ultimately, it is determined by preferences.

Of course, utility theory has had its critics, even in the early days of the Utilitarian school of thought. For example, T. Cliffe Leslie (1879), an obscure 19th century legal scholar, wrote:

There is an illusive semblance of simplicity in the Utilitarian formula . . . it assumes an unreal concord about the constituents of happiness and an unreal homogeneity of human minds in point of sensibility to different pains and pleasures. . . Nor is it possible to weigh bodily and mental pleasures and pains one against the other; no single man can pronounce with certainty about their relative intensity even for himself, far less for all his fellows. (pp. 45–46)

But even if we willingly suspend our disbelief, as most economists have done, and adopt utility theory as a useful framework for modeling *economic* decisions, *expected utility* theory still has some important limitations that several experimental studies have uncovered.

One of the earliest challenges to expected utility came from Allais (1953) and has come to be known as the "Allais paradox." Consider choosing between two alternatives, A_1 and A_2 , where

A_1 : Sure gain of \$1,000,000

A_2 : $\begin{cases} \$5,000,000 \text{ with probability } 0.10 \\ \$1,000,000 \text{ with probability } 0.89 \\ \$0 \text{ with probability } 0.01. \end{cases}$

Now, consider another two alternatives, B_1 and B_2 , where

B_1 : $\begin{cases} \$5,000,000 \text{ with probability } 0.10 \\ \$0 \text{ with probability } 0.90 \end{cases}$

B_2 : $\begin{cases} \$1,000,000 \text{ with probability } 0.11 \\ \$0 \text{ with probability } 0.89. \end{cases}$

If, like most individuals who are presented with these two binary choices, you chose A_1 and B_1 , your preferences are inconsistent with expected utility theory! To see why, observe that a preference for A_1 over A_2 implies that the expected utility of A_1 is strictly larger than that of A_2 ; hence,

$$U(1) > 0.10U(5) + 0.89U(1) + 0.01U(0) \quad (3a)$$

or

$$0.11U(1) > 0.10U(5) + 0.01U(0). \quad (3b)$$

Similarly, a preference for B_1 over B_2 implies

$$0.10U(5) + 0.90U(0) > 0.11U(1) + 0.89U(0) \quad (4a)$$

or

$$0.11U(1) < 0.10U(5) + 0.01U(0). \quad (4b)$$

But Equation 4 clearly contradicts Equation 3. To be consistent with expected utility theory, A_1 is preferred to A_2 if and only if B_2 is preferred to B_1 . The fact that many individuals in several studies have violated this preference ordering poses a serious challenge to the practical relevance of expected utility theory.¹¹

A more recent example is Kahneman and Tversky's (1979) alternative to expected utility theory called "prospect theory." They argued that individuals focus more on prospects—gains and losses—than on total wealth and that the reference point from which gains and losses are calculated can change over time. Moreover, their experiments with human subjects showed that most individuals view gains quite differently than they do losses: They are risk averse when it comes to gains and risk seeking when it comes to losses. For example, consider choosing between the following two gambles:

C_1 : Sure gain of \$240,000

C_2 : $\begin{cases} \$1,000,000 \text{ with probability } 0.25 \\ \$0 \text{ with probability } 0.75. \end{cases}$

Despite the fact that C_2 has a higher expected value than C_1 , most individuals seem to gravitate toward the sure gain, a natural display of risk aversion that can be characterized by a utility function that is concave. But now consider choosing between the following two gambles:

D_1 : Sure loss of \$750,000

D_2 : $\begin{cases} -\$1,000,000 \text{ with probability } 0.75 \\ \$0 \text{ with probability } 0.25. \end{cases}$

In this case, most individuals choose D_2 despite the fact that it is clearly a riskier alternative than D_1 . Kahneman and Tversky dubbed this behavior "loss aversion," and it can be characterized by a utility function that is convex.

This apparent asymmetry in preferences for gains and losses may not seem particularly problematic for risk management, but compare the combined outcomes of the most common choices, C_1 and D_2 , with the combined outcomes of the less popular choices, C_2 and D_1 :

$$(C_1 \text{ and } D_2) : \begin{cases} \$240,000 \text{ with probability } 0.25 \\ -\$760,000 \text{ with probability } 0.75 \end{cases}$$

$$(C_2 \text{ and } D_1) : \begin{cases} \$250,000 \text{ with probability } 0.25 \\ -\$750,000 \text{ with probability } 0.75. \end{cases}$$

C_2 and D_1 strictly dominates C_1 and D_2 ; in the former case, the gain is \$10,000 greater and the loss is \$10,000 smaller (i.e., C_2 and D_1 is equivalent to C_1 and D_2 plus a sure gain, or less loss, of \$10,000). With the choice presented in this way, and without reference to any auxiliary conditions or information, no rational individual would choose C_1 and D_2 over C_2 and D_1 . But when the two binary choices are offered separately, individuals seem to prefer the inferior choices.

Of course, one objection to this conclusion is that the test offered the two binary choices *sequentially*, not simultaneously. Although this objection is well taken, the circumstances in this example are not nearly as contrived as they might seem. For example, the London office of a multinational corporation may be faced with choices C_1 and C_2 while its Tokyo office is faced with choices D_1 and D_2 . Although locally there may not appear to be a right or wrong decision, the globally consolidated book will tell a different story. Indeed, the propensity for investors to close out winning positions too early and close out losing positions too late is well known among experienced traders—one of the first lessons one learns on a trading desk is to "cut your losses and ride your gains." The tendency for traders to increase their positions in the face of mounting losses—often called "doubling down"—is another symptom of loss aversion, one whose implications were all too real for Barings and several other financial institutions that have suffered large trading losses recently.

Another well-known challenge to expected utility is the Ellsberg (1961) paradox, in which two statistically equivalent gambles seem to be viewed very differently by the typical individual.¹² In gamble E_1 , you are asked to choose a color, red or black, after which you draw a single ball from an urn containing 100 balls, 50 red and 50 black. If you

draw a ball of your color, you receive a prize of \$10,000, otherwise you receive nothing. The terms of gamble E_2 are identical except that you draw a ball from a different urn, one containing 100 red and black balls but in unknown proportion—it may contain 100 red balls and no black balls, or 100 black balls and no red balls, or any proportion in between. What is the maximum you would pay for gamble E_1 ? And for gamble E_2 ? Alternatively, if both gambles cost the same—say, \$5,000—and you *must* choose one, which would you choose?

For most people, gamble E_2 appears to be significantly less attractive than gamble E_1 despite the fact that the probability of picking either color is identical in both gambles: 0.50. To check that the probability is indeed the same, denote by p_2 the proportion of red balls in gamble E_2 and note that p_2 can take on 101 distinct values $0/100, 1/100, \dots, 100/100$. Now, because there is no reason to favor any one proportion, the "expected" proportion can be computed by taking a weighted average of all 101 possibilities and weighting each possibility equally, which yields

$$\left(\frac{1}{101} \times \frac{0}{100}\right) + \left(\frac{1}{101} \times \frac{1}{100}\right) + \dots + \left(\frac{1}{101} \times \frac{100}{100}\right) = \frac{50}{100}.$$

Alternatively, a less formal argument is to ask what the probability could possibly be if not 50/100. In the absence of any information about the relative proportion, 50/100 is clearly the most natural hypothesis. Despite these arguments, many surveys have shown that individuals are willing to pay much less for gamble E_2 than for gamble E_1 and that when forced to choose one gamble or another at the same price, they almost always choose E_1 .

There may well be rational reasons for preferring E_1 to E_2 in other contexts, but in the simplified context in which these gambles are typically presented, it is difficult to make a compelling rationale for one or the other. This is not to say that individuals who express a preference for E_1 are irrational but, rather, that they must be incorporating other information, hypotheses, biases, or heuristics into this decision. Whether or not it is rational to include such auxiliary material in one's decision-making process depends, of course, on how relevant the material is to the specific context in which the decision is to be made. Because no single decision rule can be optimal for all circumstances, it should come as no surprise that learned responses that are nearly optimal in one context can be far from optimal in another. The value of thought experiments like the Ellsberg paradox is in illuminating certain aspects of our learned responses so that we are better able to judge their efficacy for specific purposes, such as risk management.

In particular, the Ellsberg paradox suggests that individuals have a preference regarding the uncertainty of risk. The apparent circularity of this statement (*Roget's International Thesaurus* lists risk and uncertainty as synonyms) may be resolved by recalling Knight's (1921) distinction between risk and uncertainty: Risk is the kind of randomness that can be modeled adequately by quantitative methods (e.g., mortality rates, casino gambling, equipment failure rates); the rest is uncertainty.¹³ Although Knight used this distinction to explain the seemingly disproportionate profits that accrue to entrepreneurs (they bear uncertainty, which according to Knight's theory carries a much greater reward than simply bearing risk), it also has significant implications for risk management. Indeed, the Ellsberg paradox illustrates succinctly the importance of all three P's of risk management: how much one is willing to pay for each gamble (prices), the odds of drawing red or black (probabilities), and which gamble to take and why (preferences).

Putting the Three P's Together

The challenge that lies ahead for risk-management practice is, of course, to integrate the three P's into a single and complete risk-management protocol. This daunting but essential process is a prerequisite to the growth and health of financial markets and institutions in the next century. The global financial system is becoming more complex each year, with links and interdependencies that develop and mutate day by day. Risk-management technologies must evolve in tandem.

Although the lofty goal of Total Risk Management has not yet been realized, I would like to propose two broad research agendas that show great promise for moving us closer to it. By their nature, these agendas are highly speculative, subjective, and somewhat less concrete than finished research, but the potential benefits of stimulating new ways of thinking about risk management seem well worth the hazard of making a few promises that go unfulfilled.

Preferences Revisited. The first research agenda involves revisiting the well-plowed field of preferences. Among the three P's, preferences are clearly the most fundamental and least understood aspect of risk management. Several large bodies of research have developed around these issues—in economics and finance, psychology, operations research (also known as “decision sciences”), and recently, brain and cognitive sciences. Many new insights can be gleaned from synthesizing these different strands of research into a more complete

understanding of how individuals make decisions.¹⁴ For example, are reliable methods available for measuring risk preferences quantitatively? How are risk preferences related to other aspects of personality and temperament, and can they be measured in the same ways (e.g., through surveys and psychological profiles)? What is the role of memory in determining risk-taking behavior? What can certain neurological pathologies reveal about rational decision-making capabilities and their neurophysiological origins? How do individuals learn from their own experiences and from interactions with others in economic contexts? Is it possible to construct an operational definition of rationality in the context of decision making under uncertainty? Are risk aversion and loss aversion learned traits that are acquired along the path to adulthood, or do infants exhibit these same tendencies?

Such questions lead naturally to a broadened view of economic science, one based on the principles of ecology and evolutionary biology. Unlike much of neoclassical economics and the rational expectations counterrevolution, both of which have the “look and feel” of the physical sciences, the messy empirical history of markets and economic interactions suggests a more organic interpretation. Financial markets and institutions are created, altered, and destroyed through the random and sometimes inexplicable actions of many individuals—some acting in concert, others acting independently, each acting to further her or his own goals, whatever they may be. In other words, economic systems allocate scarce resources by mutating, adapting, and evolving. In the end, economic institutions and conventions are merely another set of adaptations that evolution has given us, a metaphysical opposable thumb that has dramatically improved our chances for survival.

These ideas are not new—they owe their parentage to Edward O. Wilson's 1975 brainchild, “sociobiology”—but their application to economics and, more specifically, to financial markets has yet to be fully developed.¹⁵ If we are to understand the roots of risk preferences, it must be in the context of the survival instinct and how that has shaped economic institutions. Although this may seem too far afield to be of any practical value, recent advances in “behavioral ecology” suggest otherwise: Dynamic optimization techniques have revealed the logic of many behavioral adaptations in a variety of organisms by appealing to evolutionary principles (see, for example, Mangel and Clark 1988). Moreover, the emerging field of “evolutionary psychology”—the heir apparent to sociobiology—may also contain important insights for the origins of economic interactions.

Evolutionary psychologists have proposed compelling evolutionary arguments for a broad range of social and cultural phenomena, such as altruism, kin selection, language, mate selection, abstract thought, religion, morality, and ethics.¹⁶ Perhaps similar explanations may reveal the true nature of risk preferences and help separate those aspects that are learned from those that are inherent in our nature and nearly impossible to change. What kinds of risk preferences yield evolutionary advantages? How have evolutionary pressures influenced risk preferences? Will those pressures change over time as the nature of economic interactions changes?

But it is the recent rapprochement between evolutionary biology and molecular genetics, evidenced so eloquently by Wilson's (1994, Chapter 12) personal chronicle, that points to the most exciting and ambitious goal of all: determining the genetic basis for risk preferences. The fact that natural selection leaves its footprints in our DNA gives us a powerful tool to trace the origin of behavioral adaptations. Some progress along these lines has already been made, giving rise to a new discipline known as "behavior genetics" and populated by both cognitive scientists and molecular biologists. Using the latest techniques in DNA sequencing and computational genomics, scientists have begun to explore in earnest the heritability of behavioral traits such as anger, addiction, aggression, thrill seeking, sexual orientation, mania, depression, schizophrenia, and other aspects of temperament and personality.¹⁷

The starting point for these studies is typically a neurochemical link to certain behavioral patterns; for example, levels of the neurotransmitter dopamine in the brain seem to be correlated with thrill-seeking behavior. Once such a link is established, a genetic analysis of the corresponding neurophysiology can be conducted (e.g., identification and sequencing of the gene or genes related to dopamine receptors in the brain).¹⁸

Although the field of behavioral genetics is still in its infancy, its potential for the social sciences, and risk management in particular, is obvious. Are risk preferences simply a manifestation of a combination of other behavioral patterns, such as thrill-seeking and aggression, with different weights producing different risk tolerances, or do they have a more fundamental genetic basis? What regions of the brain are most relevant for processing risk preferences, and are these the same regions that engage in computation and quantitative reasoning? Can differences in risk preferences between two individuals be determined through genetic comparisons, and if so, what might the implications be for risk management, both private and social?

Risk in Broader Contexts. The second research agenda is motivated by the fact that risk is a common feature of many human endeavors; hence, much can be gained from considering how other disciplines deal with risk measurement and management. For example, risk assessment is an integral component of chemical, aeronautical, astronautical, and nuclear engineering, epidemiology and public health policy, biomedical technology, and the insurance industries. In each of these fields, academic research is intimately tied to industry applications, which yields practical risk-management policies that may contain novel insights for financial risk management. And recent innovations in financial risk management may provide new ways of thinking about risk in nonfinancial contexts. In either case, it is clear that risk is a universal phenomenon and may be better understood by studying it in a broad framework.

Such a framework is hinted at in the influential work of sociologist Charles Perrow (1984), in which he argued that certain catastrophes are unavoidable consequences of systems that are simply too complex and too unforgiving. He described in great detail the pathologies of the Three Mile Island nuclear reactor breach, aircraft and air-traffic-control accidents, various petrochemical plant explosions, and a host of other man-made disasters, and he made a compelling case that these accidents are not pathological at all but are "normal" for organizations of such complexity. By identifying specific organizational features that are likely to generate "normal accidents," Perrow provided useful guidelines for thinking about risk management in a broad context. In particular, he categorized systems along two dimensions—the degree to which the individual components can interact with each other and the reliance of one component's functionality on another's. Systems in which individual components can interact in complex ways (systems that exhibit "interactive complexity") and in which the functions of many components are highly dependent on those of other components (systems that exhibit "tight coupling") are prime candidates for normal accidents.

These ideas, and the industrial accidents that inspired them, have had a significant impact on the way industries and policymakers view risks, and they are responsible for at least two new journals and a burgeoning literature on "high reliability organizations" and the management of enterprise-wide risks.¹⁹ Although much of this literature is descriptive and qualitative, its relevance for financial risk management is clear: Accidents are normal in industrial systems so complex and nonlinear that small and unpredictable errors in

human judgment can often cascade quickly and inexorably into major catastrophes. The challenge is to quantify the notions of interactive complexity and tight coupling so that intelligent trade-offs between risk and reward can be properly made, in both financial and nonfinancial contexts. Perhaps the new mathematics of “nonlinear dynamical systems”—deterministic nonlinear equations that exhibit extraordinarily complex behavior—can play a role in defining these trade-offs.

The Future of Risk Management

If the two research agendas outlined earlier seem too far removed from the daily focus of risk-management practices, consider the fact that the centerpiece of each of the most prominent failures of financial risk-management systems in the past few years—Procter & Gamble, Gibson Greetings, Orange County, and Barings—is human judgment and risk preferences. Alternatively, Street-smart traders often attribute the ebb and flow of financial fortunes to just two factors: fear and greed. Although connecting these aspects of human behavior with biology may require a stretch of the imagination, the distance is shrinking day by day.

Consilience. The fact that the two research agendas proposed cut across so many different disciplines—economics and finance, statistics, biology, and the brain and cognitive sciences—may well be part of a growing trend, a manifestation of Wilson’s (1998) notion of consilience: “literally a ‘jumping together’ of knowledge by the linking of facts and fact-based theory across disciplines to create a common groundwork of explanation” (p. 8). In considering the state of the social sciences, Wilson wrote:

The full understanding of utility will come from biology and psychology by reduction to the elements of human behavior followed by bottom-up synthesis, not from the social sciences by top-down inference and guesswork based on intuitive knowledge. It is in biology and psychology that economists and other social scientists will find the premises needed to fashion more predictive models, just as it was in physics and chemistry that researchers found premises that upgraded biology. (p. 206)

If financial economics is to graduate to the level of a true scientific discipline, a promising starting point might be the sociobiological foundations of the three P’s of risk management.

A TRM Protocol. Despite the fact that the two research agendas outlined here contain a series of concrete issues to be investigated, it is easy to lose

sight of the ultimate goal of a fully integrated TRM protocol. What would such a protocol look like upon completion of the proposed research?

A TRM protocol for an institution might consist of the following five phases. The first phase is an analysis of the organization’s structure to determine its susceptibility to normal accidents (i.e., a quantitative analysis of its interactive complexity and tightness of coupling). Such an analysis can be performed without reference to any of the three P’s because the focus is on the *system* and the limitations embedded in its structure, not on the likelihood or impact of encountering such limitations.

The second phase—probabilities—is a risk-assessment process in which the probabilities of various events and scenarios are either postulated or estimated. The distinction between objective and subjective probabilities should be clarified at this stage, and all probabilities should be checked for mutual consistency. Preferences and prices might also play a role here to the extent that they can be used or restricted in some fashion to estimate probabilities more accurately (see, for example, Shimko 1993, Rubinstein 1994, Jackwerth and Rubinstein, and Ait-Sahalia and Lo 1998a).

The third phase—prices—involves determining the economic consequences of various events and scenarios, either by using market prices or by computing equilibrium prices (which would require preferences and probabilities) for nonmarketed or illiquid instruments.

The fourth phase—preferences—consists of a comprehensive risk-attitudes inventory of all the relevant decision makers and a determination of the overall business objectives of the enterprise. Individual preferences can be determined through several means: psychological and risk profiles (questionnaires), historical performance records, and perhaps even physiological (blood levels of testosterone and cortisol) and genetic analysis (genetic predisposition for risk-processing abilities).²⁰ Once the major decision makers’ risk preferences and the corporate objectives have been determined, it will be possible to analyze risk preferences in light of various compensation structures to check that the possible interactions are consistent with those objectives. For example, if an individual is risk neutral and his or her compensation consists primarily of warrants on the company’s stock, the individual’s behavior might not be consistent with the maximization of shareholder wealth.²¹ Such considerations could be used not only to redesign compensation packages but also to screen for employees with risk preferences consistent with existing compensation structures and corporate objectives.

And the fifth and final phase involves the development and implementation of an automated, real-time risk-monitoring system that can keep track of any significant changes in the three P's, including changes in key decision makers' compensation levels and, consequently, their wealth (which might affect their preferences), changes in institutional structure, and changes in business conditions. Although this might seem out of reach today, recent advances in expert systems, natural language processing, computational learning algorithms, and computing power might allow us to build such systems in the not-too-distant future.

Such a TRM protocol can also be easily adapted to an individual's decision-making process, and this might be the most important application of all. Because of the shift from defined-benefit to defined-contribution pension plans in the majority of corporations today, individuals are being charged with the awesome responsibility of planning for their own retirement. If we can truly integrate prices, probabilities, and preferences in a framework that

enables individuals and institutions to manage their respective risks systematically and successfully, we will have achieved the ultimate Utilitarian mandate: the greatest good for the greatest number.

This research was partially supported by the MIT Laboratory for Financial Engineering, the National Science Foundation, and the PricewaterhouseCoopers Risk Institute. I am grateful to Samantha Arrington, Wesley Chan, John Cox, Charles Harris, Martin Haugh, Leonid Kogan, Ken Kosik, Nancy Lo, Hal Lux, Charles Perrow, Steve Ross, Ross Stevens, Jiang Wang, the seminar participants at the Bond Analysts Society of Boston, the International Swap Dealers Association 13th Annual General Meeting, Mobil Corporation, the New York Society of Quantitative Analysts, and the Wharton Risk Measurement and Management Conference for valuable comments and discussion.

Notes

1. See, for example, Bernstein's (1996) lively historical account of risk.
2. I thank Zvi Bodie for suggesting this term.
3. For example, the price, P_t , of any financial security that pays a stream of dividends D_{t+1}, D_{t+2}, \dots must satisfy the following relationship:

$$P_t = E_t \left[\sum_{\tau=1}^{\infty} \frac{U'_{t+\tau}(C_{t+\tau})}{U'_t(C_t)} D_{t+\tau} \right],$$

where $U'_t(C_t)$ and $U'_{t+\tau}(C_{t+\tau})$ are the marginal utilities of consumption at dates t and $t + \tau$, respectively.

4. My colleague Jiang Wang has observed that the term "asset pricing" implies an inordinate focus on prices, often to the exclusion of other interesting and, in some cases, equally important economic phenomena (e.g., quantities). Perhaps this focus is another manifestation of this article's theme: Prices alone cannot provide a complete understanding of the nature of financial risks and rewards; other aspects of market interactions—probabilities and preferences—must be considered. Wang has suggested a simple but compelling alternative to asset pricing: "asset markets" (as in "asset-market models" instead of "asset-pricing models").
5. In addition to Merton's seminal paper (1973a), several other important contributions to the finance literature are responsible for our current understanding of dynamic spanning. In particular, see Cox and Ross (1976), Harrison and Kreps (1979), Huang (1985a, 1985b, 1987), and Duffie and Huang (1985).
6. Although Merton (1973a) rederived the Black-Scholes formula using arbitrage arguments alone, he was able to do so only because of his use of continuous-time stochastic processes. The links among continuous-time models, arbitrage, and equilibrium are complex and have given rise to a large and still-growing literature now known as mathematical

finance. See Harrison and Kreps, Duffie and Huang, and Merton (1992) for further discussion.

7. See, for example, Hald (1990, Chapter 3). Also, Bernstein (1996) and Sherden (1998) provide very entertaining and informative accounts of the checkered history of probability, risk, and forecasting.
8. This surely must be one of the earliest examples of the use of a financial principle—the absence of arbitrage—to support a mathematical proposition!
9. See also Samuelson (1983), which is an expanded version of his *tour de force* that includes an excellent discussion in Appendix C of more-recent developments (as of 1983) in expected utility theory, mean-variance analysis, and general portfolio theory. And for a fascinating account of the origins of *Foundations*, see Samuelson (1998).
10. Briefly, if \succeq denotes a preference relationship (that is, $A \succeq B$ means A is preferred or indifferent to B), then the following axioms are sufficient for expected utility theory to hold:
 - Completeness.* For any two gambles A and B , either $A \succeq B$ or $B \succeq A$ or both.
 - Transitivity.* If $A \succeq B$ and $B \succeq C$, then $A \succeq C$.
 - Continuity.* If $A \succ B \succeq C$, then there exists some $\lambda \in [0, 1]$ such that B is not different in preference from $\lambda A + (1 - \lambda)C$.
 - Independence.* For any two gambles A and B , $A \succeq B$ if and only if $\lambda A + (1 - \lambda)C \succeq \lambda B + (1 - \lambda)C$ for all C and $\lambda \in [0, 1]$. Herstein and Milnor (1953) provide a rigorous treatment of von Neumann and Morgenstern's derivation. See Fishburn (1970) and Kreps (1988) for a thorough modern exposition of expected utility.
11. See, for example, Morrison (1967), Raiffa (1968), Moskowitz (1974), and Slovic and Tversky (1974).
12. The E_1, E_2 example is a slightly modified version of Ellsberg's original thought experiment, modified to simplify the exposition.

13. In fact, Ellsberg (p. 653) acknowledged that Knight proposed the same thought experiment of an individual choosing between two urns, one with a known proportion of red and black balls and another with an unknown proportion.
14. Simon's 1982 contributions to this literature are still remarkably timely, and their implications have yet to be fully explored. For more recent contributions, see Kahneman, Slovic, and Tversky (1982); Hogarth and Reder (1986); Gigerenzer and Murray (1987); Dawes (1988); Fishburn (1988); Keeney and Raiffa (1993); Plous (1993); Sargent (1993); Thaler (1993); Damasio (1994); Arrow et al. (1996); Picard (1997); Pinker (1997); and Rubinstein (1998).
15. Students of the history of economic thought will no doubt recall that Thomas Malthus used biological arguments—the fact that populations increase at geometric rates whereas natural resources increase only arithmetically—to draw economic implications and that both Darwin and Wallace were influenced by these arguments (see Hirshleifer 1977 for further details). Also, Joseph Schumpeter's views of business cycles, entrepreneurs, and capitalism have an evolutionary flavor to them; in fact, his notions of "creative destruction" and "bursts" of entrepreneurial activity are similar in spirit to natural selection and punctuated equilibria. Recently, economists and biologists have begun to explore these connections in several directions: direct extensions of sociobiology to economics (Becker 1976; Hirshleifer; Tullock 1979), evolutionary game theory (Smith 1982; Weibull 1995), evolutionary economics (Nelson and Winter 1982; Andersen 1994; Englund 1994), and economics as a complex system (Anderson, Arrow, and Pines 1988). See Hodgson (1995) for a collection of studies on economics and biology.
16. See Pinker (1993, 1997), Barkow, Cosmides, and Tooby (1992), and Crawford and Krebs (1998).
17. See Hamer and Copeland (1998) for an excellent and up-to-date survey of behavioral genetics. Other recent surveys include Plomin (1990), Steen (1996), Barondes (1998), and Wright (1998). Skeptics might argue that the entire field of behavioral genetics rests on one side of the age-old nature-versus-nurture debate (for a recent study that weighs in on the other side, see Harris 1995, 1998). As research progresses in both genetics and psychology, however, this debate is becoming less heated in some respects. Nature and nurture seem to work together to determine behavior, and the more relevant question is: How?
18. See Benjamin et al. (1996) for the specific example of thrill seeking and dopamine receptors. Plomin, Owen, and McGuffin (1994) provide a detailed survey of this burgeoning literature.
19. See, in particular, the *Industrial Crisis Quarterly*, the *Journal of Contingencies and Crisis Management*, and Sagan (1993), Perrow (1994), La Porte (1996), Rochlin (1996), and Reason (1997).
20. These last two possibilities are no doubt the most controversial, and they raise a number of challenging issues regarding individual privacy, social policy, and ethics. Such issues are not new, but they have received even greater attention in the wake of recent breakthroughs in biotechnology (see, for example, Weiss and Straughan 1996). Although a simple resolution of these issues in the near future may be too much to hope for, the sheer volume of biotechnology applications currently being developed will require that clear guidelines be established soon.
21. In particular, he will have incentives to take on more risk, in some cases, even at the expense of corporate profits.

References

- Ait-Sahalia, Y., and A. Lo. 1998a. "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices." *Journal of Finance*, vol. 53, no. 2 (June):499–548.
- . 1998b. "Nonparametric Risk Management and Implied Risk Aversion." *Journal of Econometrics* (forthcoming).
- Allais, M. 1953. "Fondements d'une Théorie Positive des Choix Comportant un Risque et Critique des Postulats et Axiomes de l'École Américaine." *Econometrica*, vol. 21, no. 5: 503–546. Translated into English by M. Allais and O. Hagen, eds. 1979. *Expected Utility Hypotheses and the Allais Paradox*. Dordrecht, Netherlands: D. Reidel.
- Andersen, E. 1994. *Evolutionary Economics: Post-Schumpeterian Contributions*. London: Pinter.
- Anderson, P., K. Arrow, and D. Pines, eds. 1988. *The Economy as an Evolving Complex System*. Reading, MA: Addison-Wesley Publishing.
- Arrow, K., E. Colombatto, M. Perlman, and C. Schmidt. 1996. *The Rational Foundations of Economic Behaviour*. New York: St. Martin's Press.
- Barkow, J., L. Cosmides, and J. Tooby. 1992. *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. Oxford, U.K.: Oxford University Press.
- Barondes, S. 1998. *Mood Genes: Hunting for Origins of Mania and Depression*. New York: W.H. Freeman.
- Becker, G. 1976. "Altruism, Egoism, and Genetic Fitness: Economics and Sociobiology." *Journal of Economic Literature*, vol. 14, no. 3:817–826.
- Benjamin, J., L. Li, C. Patterson, B. Greenberg, D. Murphy, and D. Hamer. 1996. "Population and Familial Association between the D4 Dopamine Receptor Gene and Measures of Novelty Seeking." *Nature Genetics*, vol. 12, no. 1:81–84.
- Bernoulli, D. 1738. "Specimen Theoriae Novae de Mensura Sortis." *Commentarii Academiae Scientiarum Imperialis Petropolitanae*, vol. 5, no. 2:175–192. Translated into English by L. Sommer. 1954. "Exposition of a New Theory on the Measurement of Risk." *Econometrica*, vol. 22, no. 1:23–36.
- Bernstein, P. 1996. *Against the Gods: The Remarkable Story of Risk*. New York: John Wiley & Sons.
- Bick, A. 1990. "On Viable Diffusion Price Processes of the Market Portfolio." *Journal of Finance*, vol. 45, no. 2 (June):673–689.
- Black, F., and M. Scholes. 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, vol. 81, no. 3:637–659.
- Breeden, D. 1979. "An Intertemporal Capital Pricing Model with Stochastic Investment Opportunities." *Journal of Financial Economics*, vol. 7, no. 3:265–296.
- Cox, J., and S. Ross. 1976. "The Valuation of Options for Alternative Stochastic Processes." *Journal of Financial Economics*, vol. 3, nos. 1–2:145–166.
- Cox, J., J. Ingersoll, and S. Ross. 1985. "An Intertemporal General Equilibrium Model of Asset Prices." *Econometrica*, vol. 53, no. 2:363–384.

- Crawford, C., and D. Krebs, eds. 1998. *Handbook of Evolutionary Psychology: Ideas, Issues, and Applications*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Damasio, A. 1994. *Descartes' Error: Emotion, Reason, and the Human Brain*. New York: Avon Books.
- Dawes, R. 1988. *Rational Choice in an Uncertain World*. New York: Harcourt Brace.
- De Finetti, B. 1937. "La Prévision: Ses Lois Logiques, Ses Sources Subjectives." *Annales de l'Institut Henri Poincaré*, vol. 7, no. 1:1–68. Translated into English in H. Kyburg and H. Smokler, eds. 1964. *Studies in Subjective Probability*. New York: John Wiley & Sons.
- Duffie, D., and C. Huang. 1985. "Implementing Arrow–Debreu Equilibria by Continuous Trading of Few Long-Lived Securities." *Econometrica*, vol. 53, no. 6:1337–56.
- Ellsberg, D. 1961. "Risk, Ambiguity, and the Savage Axioms." *Quarterly Journal of Economics*, vol. 75, no. 4:643–669.
- Englund, R., ed. 1994. *Evolutionary Concepts in Contemporary Economics*. Ann Arbor, MI: University of Michigan Press.
- Fishburn, P. 1970. *Utility Theory for Decision Making*. New York: John Wiley & Sons.
- . 1988. *Nonlinear Preference and Utility Theory*. Baltimore, MD: Johns Hopkins University Press.
- Gigerenzer, G., and D. Murray. 1987. *Cognition as Intuitive Statistics*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Hald, A. 1990. *A History of Probability and Statistics and Their Applications before 1750*. New York: John Wiley & Sons.
- Hamer, D., and P. Copeland. 1998. *Living with Our Genes*. New York: Doubleday.
- Harris, J. 1995. "Where Is the Child's Environment? A Group Socialization Theory of Development." *Psychological Review*, vol. 102, no. 3:458–489.
- . 1998. *The Nurture Assumption: Why Children Turn Out the Way They Do*. New York: The Free Press.
- Harrison, M., and D. Kreps. 1979. "Martingales and Arbitrage in Multiperiod Securities Markets." *Journal of Economic Theory*, vol. 2, no. 3:381–408.
- He, H., and H. Leland. 1993. "On Equilibrium Asset Price Processes." *Review of Financial Studies*, vol. 6, no. 3 (Fall):593–617.
- Herstein, I., and J. Milnor. 1953. "An Axiomatic Approach to Measurable Utility." *Econometrica*, vol. 21, no. 2:291–297.
- Hirshleifer, J. 1977. "Economics from a Biological Viewpoint." *Journal of Law and Economics*, vol. 20, no. 1:1–52.
- Hodgson, G., ed. 1995. *Economics and Biology*. Cheltenham, U.K.: Edward Elgar Publishing.
- Hogarth, R., and M. Reder. 1986. *Rational Choice: The Contrast between Economics and Psychology*. Chicago, IL: University of Chicago Press.
- Huang, C. 1985a. "Information Structure and Equilibrium Asset Prices." *Journal of Economic Theory*, vol. 35, no. 1:33–71.
- . 1985b. "Information Structures and Viable Price Systems." *Journal of Mathematical Economics*, vol. 14, no. 3:215–240.
- . 1987. "An Intertemporal General Equilibrium Asset Pricing Model: The Case of Diffusion Information." *Econometrica*, vol. 55, no. 1:117–142.
- Jackwerth, J. 1998. "Recovering Risk Aversion from Option Prices and Realized Returns." Working paper. Haas School of Business, University of California, Berkeley.
- Jackwerth, J., and M. Rubinstein. 1996. "Recovering Probability Distributions from Contemporary Security Prices." *Journal of Finance*, vol. 51, no. 5 (December):1611–31.
- Kahneman, D., P. Slovic, and A. Tversky. 1982. *Judgment under Uncertainty: Heuristics and Biases*. Cambridge, U.K.: Cambridge University Press.
- Kahneman, D., and A. Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica*, vol. 47, no. 2:263–291.
- Keeney, R., and H. Raiffa. 1993. *Decisions with Multiple Objectives: Preferences and Value Trade-Offs*. Cambridge, U.K.: Cambridge University Press.
- Knight, F. 1921. *Risk, Uncertainty, and Profit*. Boston, MA: Houghton Mifflin.
- Kreps, D. 1988. *Notes on the Theory of Choice*. Boulder, CO: Westview Press.
- La Porte, T. 1996. "High Reliability Organizations: Unlikely, Demanding and at Risk." *Journal of Contingencies and Crisis Management*, vol. 4, no. 2:60–72.
- Leslie, T. Cliffe. 1879. *Essays in Political and Moral Philosophy*. London: Longmans, Green.
- Lo, A., and J. Wang. 1995. "Implementing Option Pricing Models when Asset Returns Are Predictable." *Journal of Finance*, vol. 50, no. 1 (March):87–129.
- Lucas, R. 1978. "Asset Prices in an Exchange Economy." *Econometrica*, vol. 46, no. 6:1429–46.
- Mangel, M., and C. Clark. 1988. *Dynamic Modeling in Behavioral Ecology*. Princeton, NJ: Princeton University Press.
- Merton, R. 1973a. "Rational Theory of Option Pricing." *Bell Journal of Economics and Management Science*, vol. 4, no. 1:141–183.
- . 1973b. "An Intertemporal Capital Asset Pricing Model." *Econometrica*, vol. 41, no. 5:867–887.
- . 1992. *Continuous-Time Finance*. Revised ed. Oxford, U.K.: Basil Blackwell.
- Morgan Guaranty Trust Company. 1995. *Introduction to RiskMetrics*. 4th ed. New York: Morgan Guaranty Trust.
- Morrison, D. 1967. "On the Consistency of Preferences in Allais' Paradox." *Behavioral Science*, vol. 12, no. 3:373–383.
- Moskowitz, H. 1974. "Effects of Problem Representation and Feedback on Rational Behavior in Allais and Morlat-Type Problems." *Decision Sciences*, vol. 5, no. 2:225–242.
- Nelson, R., and S. Winter. 1982. *An Evolutionary Theory of Economic Change*. Cambridge, MA: Belknap Press of Harvard University Press.
- Perrow, C. 1984. *Normal Accidents: Living with High Risk Technologies*. New York: Basic Books.
- . 1994. "The Limits of Safety: The Enhancement of a Theory of Accidents." *Journal of Contingencies and Crisis Management*, vol. 2, no. 4:212–220.
- Picard, R. 1997. *Affective Computing*. Cambridge, MA: MIT Press.
- Pinker, S. 1993. *The Language Instinct*. New York: Morrow.
- . 1997. *How the Mind Works*. New York: W.W. Norton.
- Plomin, R. 1990. *Nature and Nurture: An Introduction to Human Behavioral Genetics*. Pacific Grove, CA: Brooks/Cole Publishing.
- Plomin, R., M. Owen, and P. McGuffin. 1994. "The Genetic Basis of Complex Human Behaviors." *Science*, vol. 264, no. 5166:1733–39.
- Plous, S. 1993. *The Psychology of Judgment and Decision Making*. New York: McGraw-Hill.
- Raiffa, H. 1968. *Decision Analysis: Introductory Lectures on Choice under Uncertainty*. Reading, MA: Addison Wesley.
- Ramsey, F. 1926. "Truth and Probability." *Foundations of Mathematics and Other Logical Essays*. Edited by R. Braithwaite. New York: Harcourt Brace.

- Reason, J. 1997. *Managing the Risks of Organizational Accidents*. Hants, U.K.: Ashgate.
- Rochlin, G. 1996. "Reliable Organizations: Present Research and Future Directions." *Journal of Contingencies and Crisis Management*, vol. 4, no. 2:55–59.
- Rubinstein, A. 1998. *Modeling Bounded Rationality*. Cambridge, MA: MIT Press.
- Rubinstein, M. 1994. "Implied Binomial Trees." *Journal of Finance*, vol. 49, no. 3 (July):771–818.
- Sagan, S. 1993. *The Limits of Safety: Organizations, Accidents, and Nuclear Weapons*. Princeton, NJ: Princeton University Press.
- Samuelson, P. 1947. *Foundations of Economic Analysis*. Cambridge, MA: Harvard University Press.
- . 1983. *Foundations of Economic Analysis*. Enlarged edition. Cambridge, MA: Harvard University Press.
- . 1998. "How Foundations Came to Be." *Journal of Economic Literature*, vol. 36, no. 3:1375–86.
- Sargent, T. 1993. *Bounded Rationality in Macroeconomics*. Oxford, U.K.: Clarendon Press.
- Savage, L. 1954. *Foundations of Statistics*. New York: John Wiley & Sons.
- Sharpe, W. 1964. "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk." *Journal of Finance*, vol. 19, no. 3 (July): 425–442.
- Sherden, W. 1998. *The Fortune Sellers: The Big Business of Buying and Selling Predictions*. New York: John Wiley & Sons.
- Shimko, D. 1993. "Bounds of Probability." *Risk*, vol. 6, no. 1:33–37.
- Simon, H. 1982. *Models of Bounded Rationality*, vols. 1 and 2. Cambridge, MA: MIT Press.
- Slovic, P., and A. Tversky. 1974. "Who Accepts Savage's Axioms?" *Behavioral Science*, vol. 19, no. 2:368–373.
- Smith, A. 1776. *An Inquiry into the Nature and Causes of the Wealth of Nations*. Reprinted 1976. Oxford, U.K.: Clarendon Press.
- Smith, J. Maynard. 1982. *Evolution and the Theory of Games*. Cambridge, U.K.: Cambridge University Press.
- Steen, G. 1996. *DNA and Destiny: Nature and Nurture in Human Behavior*. New York: Plenum Press.
- Thaler, R., ed. 1993. *Advances in Behavioral Finance*. New York: Russell Sage Foundation.
- Tullock, G. 1979. "Sociobiology and Economics." *Atlantic Economic Journal*, vol. 7, no. 1:1–10.
- von Neumann, J., and O. Morganstern. 1944. *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.
- Weibull, J. 1995. *Evolutionary Game Theory*. Cambridge, MA: MIT Press.
- Weiss, M., and R. Straughan. 1996. *Improving Nature? The Science and Ethics of Genetic Engineering*. Cambridge, U.K.: Cambridge University Press.
- Wilson, E. 1975. *Sociobiology: The New Synthesis*. Cambridge, MA: Belknap Press of Harvard University Press.
- . 1994. *Naturalist*. Washington, DC: Island Press.
- . 1998. *Consilience: The Unity of Knowledge*. New York: Alfred A. Knopf.
- Wright, W. 1998. *Born That Way: Genes, Behavior, Personality*. New York: Alfred A. Knopf.