



2015

Risk, Uncertainty, and "Super-Risk"

José Luis Bermúdez

Michael S. Pardo

Follow this and additional works at: <http://scholarship.law.nd.edu/ndjlepp>



Part of the [Ethics and Professional Responsibility Commons](#)

Recommended Citation

29 Notre Dame J.L. Ethics & Pub. Pol'y 471 (2015)

This Article is brought to you for free and open access by the Notre Dame Journal of Law, Ethics & Public Policy at NDLScholarship. It has been accepted for inclusion in Notre Dame Journal of Law, Ethics & Public Policy by an authorized administrator of NDLScholarship. For more information, please contact lawdr@nd.edu.

Risk, Uncertainty, and "Super-Risk"

Cover Page Footnote

* Associate Provost for Strategic Planning and Professor of Philosophy, Texas A&M University. ** Henry Upson Sims Professor of Law, University of Alabama School of Law. For helpful comments on previous drafts, we thank Ron Allen, Ed Cheng, Meredith Render, Alex Stein, and Fred Vars. Prof. Pardo thanks Dean Mark Brandon and the University of Alabama Law School Foundation for generous research support.

RISK, UNCERTAINTY, AND “SUPER-RISK”

JOSÉ LUIS BERMÚDEZ* & MICHAEL S. PARDO**

ABSTRACT

Risk is a pervasive feature of law and public policy. Decision-making in these domains often takes place in the absence of certainty and with awareness that errors may be made and predictions may fail. Within law—as within the social and physical sciences, medicine, economics, finance, and countless other domains—a primary focus of practical and scholarly inquiries is the extent to which risks can be measured and managed. In each of these domains, risk analysis typically employs the basic tools of decision theory (probability and utility) to measure the likelihood as well as the costs and benefits associated with possible outcomes. Risk analysis also often makes use of the familiar (but confusing) distinction between decisions made in conditions of “risk” (roughly, the relevant likelihoods and costs are quantifiable) and decisions made in conditions of “uncertainty” (roughly, the possibilities are either unknown or not amenable to quantification).

Beginning with the risk-uncertainty distinction, but altering its terminology, we argue that there is a fundamentally important type of risk that has been systematically ignored. We call it “super-risk.” Super-risk occurs when, at the time of decision, decision-makers believe they are in conditions of risk (what we call “actuarial decision-making”), but they do not know whether they are in an actuarial or an uncertain environment. Super-risk gives rise to a particular type of inferential problem, with significant practical consequences, when decision-makers proceed under the assumption that they are in an actuarial environment but they are in fact in an uncertain one. Super-risk has the potential to arise in any decision-making domain with uncertain outcomes, but it is more prone to arise with decision-making in domains such as law, public policy, economics, finance, and the social sciences rather than in domains such as the physical sciences, medicine, and insurance. Our goal in this Article is to introduce the general idea of super-risk and to explain its features and sources.

INTRODUCTION

Risk is pervasive in law and public policy. To illustrate this risk, consider two seemingly disparate examples: financial regulation and trials. Decision-making in each of these domains is made in conditions of

* Associate Provost for Strategic Planning and Professor of Philosophy, Texas A&M University.

** Henry Upson Sims Professor of Law, University of Alabama School of Law. For helpful comments on previous drafts, we thank Ron Allen, Ed Cheng, Meredith Render, Alex Stein, and Fred Vars. Prof. Pardo thanks Dean Mark Brandon and the University of Alabama Law School Foundation for generous research support.

risk, and legal rules promulgated in both domains focus on the appropriate levels and distribution of this risk. For instance, capital requirements on banks are a type of financial regulation designed to manage risk.¹ Similarly, burdens and standards of proof are used at trial to protect against “the risk of error inherent in the truth-finding process.”² In addition to these specific examples, the law generally regulates decision-making by individuals, corporations, industries, and the government as they operate in uncertain environments. Environmental regulations, food and drug safety, tort law, criminal law, property law, contract law, and constitutional law all involve decision-making under such conditions.³

In most domains—including law, public policy, the physical and social sciences, medicine, and technology—a standard theoretical framework is used to understand, measure, and manage risk.⁴ This framework employs probabilistic models and statistical information to predict possible outcomes. The basic conception of risk underlying

1. See Charles K. Whitehead, *Reframing Financial Regulation*, 90 B.U. L. REV. 1, 25 (2010) (“For banks, regulatory capital [requirements] cushion[] against the risk of loss from a portfolio of loans, . . .” among other things.).

2. The Supreme Court of the United States has explained that burdens of proof and other procedural rules regulating the process of legal proof “are shaped by the risk of error inherent in the truth-finding process[.]” *Santosky v. Kramer*, 455 U.S. 745, 757 (1982) (internal quotation omitted).

3. See, e.g., Daniel Farber, *Symposium Introduction: Navigating the Intersection of Environmental Law and Disaster Law*, 2011 BYU L. REV. 1783, 1792 (2011) (“Much of environmental law involves principles for determining the seriousness of risk and the extent to which society should invest in reducing those risks.”); Matthew D. Adler, *Against “Individualized Risk”: A Sympathetic Critique of Risk Assessment*, 153 U. PA. L. REV. 1121, 1164–69 (2005) (discussing risk assessment by the Food and Drug Administration); Richard A. Posner, *A Theory of Negligence*, 1 J. LEG. STUD. 29, 33 (1972) (conceptualizing negligence in terms of the costs associated with preventing expected harm); Jonathan Remy Nash, *The Supreme Court and the Regulation of Risk in Criminal Law Enforcement*, 92 B.U. L. REV. 171, 172 (2012) (“Criminal law provides a natural home for risk regulation.”); Henry E. Smith, *Property and Property Rules*, 79 N.Y.U. L. REV. 1719, 1724 (2004) (“The problem of information in protecting property is characterized not just by risk but also by varying degrees of uncertainty. First-order production of information on the relevant asset or class of assets involves not only managing risk, but also turning uncertainty into more manageable risk.”); Alan Schwartz & Robert E. Scott, *Contract Theory and the Limits of Contract Law*, 113 YALE L. J. 541, 550 (2003) (advancing a wealth-maximization theory of contract law on the ground that firms are “risk neutral” and seek to maximize expected profits); Adrian Vermeule, *Precautionary Principles in Constitutional Law*, 4 J. LEG. ANALYSIS 181, 181 (2012) (applying “the insights of risk analysis” from disciplines such as “decision theory and game theory, welfare economics, political science, and psychology” to “constitutional law and theory”).

4. See, e.g., L. Jonathan Cohen, *Bayesian versus Baconianism in the Evaluation of Medical Diagnoses*, 31 BRIT. J. PHI. SCI. 45 (1980) (“Medicine is an excellent example of a field in which it is often necessary to make inferences under conditions of uncertainty, either because further delays in therapy would be too risky or because further tests or examinations would be too costly.”); JUDEA PEARL, *CAUSALITY: MODELS, REASONING, AND INFERENCE* ch. 3 (2d ed. 2009) (discussing statistical models in the social sciences and economics); MICHAEL A. BISHOP & J.D. TROUT, *EPISTEMOLOGY AND THE PSYCHOLOGY OF HUMAN JUDGMENT* 24–53 (2005) (discussing statistical prediction rules in different domains, including the social sciences, medicine, and law); Robyn M. Dawes et al., *Clinical versus Actuarial Judgment*, 243 SCIENCE 1668 (1989) (advocating for the use of probabilistic models of decision-making in any domain for which empirically established relations are available).

these models involves three components: (1) possible outcomes or states of the world; (2) costs or benefits associated with each outcome or state; and (3) probabilities or likelihoods that each outcome or state obtains or will obtain.⁵

This basic conception of risk is the starting point of this Article, which argues that there is a conceptually distinct type of risk, called "super-risk," which is fundamentally important and systematically ignored. The goals of this Article are to introduce the general concept of super-risk and to explain its features and sources. For reasons explored in-depth in this Article, super-risk has the potential to arise in decision-making in *any* of the aforementioned domains; however, it is more likely to arise in the domains of law, public policy, economics, finance, and the social sciences, than in the domains of medicine, the physical sciences, and insurance. The concept of super-risk is illustrated by detailed examples from law and finance. Thereafter, decision-making in the contexts of law and finance is contrasted with decision-making in the contexts of science, medicine, and insurance. All of the examples discussed herein are intended to illuminate the more general phenomenon and its features across all domains.

In order to properly understand the concept of super-risk, one must first understand Frank Knight's familiar, but confusing distinction between decisions made under conditions of "risk" and decisions made under conditions of "uncertainty."⁶ Decisions under *risk* involve circumstances in which the probabilities and costs associated with possible outcomes can be quantified; whereas decisions under *uncertainty* involve circumstances in which the probabilities and costs associated with possible outcomes are not amenable to quantification.⁷ *Super-risk* arises when, at the time of a decision, the decision-maker believes she is making decisions in conditions of risk, when she is actually in conditions of uncertainty. Furthermore, this scenario creates the possibility for a particular type of inferential error.⁸

The first context in which the Article illustrates super-risk—financial decision-making—is discussed against the backdrop of recent financial crises. Scholars have identified various causes and conditions giving rise to these crises.⁹ This Article argues that super-risk was also a

5. The standard conception of risk is based on the theory of expected utility. *See generally* RICHARD C. JEFFREY, *THE LOGIC OF DECISION* (2d ed. 1983); R. DUNCAN LUCE & HOWARD RAIFFA, *GAMES AND DECISIONS: INTRODUCTION AND CRITICAL SURVEY* (1957). For an overview, *see* JOSÉ LUIS BERMÚDEZ, *DECISION THEORY AND RATIONALITY* (2009). Decision theorists and economists have proposed alternatives to the expected utility theory. *See, e.g.*, Chris Starmer, *Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice Under Risk*, 38 J. ECON. LIT. 332 (2000); Robert Sugden, *Why Be Consistent? A Critical Analysis of Consistency Requirements in Choice Theory*, 52 *ECONOMICA* 167 (1985).

6. FRANK H. KNIGHT, *Preface to the First Edition of RISK, UNCERTAINTY, AND PROFIT* 19–20 (5th impression 1940).

7. *Id.*

8. As we will explain, super-risk is a special case of "epistemic risk," or the risk that one's assessment of risk is faulty. *See* Peter Gärdenfors & Nils-Eric Sahlin, *Unreliable Probabilities, Risk Bearing, and Decision Making*, 33 *SYNTHESE* 361 (1982).

9. *See, e.g.*, RICHARD A. POSNER, *THE CRISIS OF CAPITALIST DEMOCRACY* (2010).

cause of the financial crises. More specifically, this Article contends that super-risk was one of the root causes of recent market turmoil and the investment catastrophe.

The second context in which this Article illustrates super-risk—fact-finding at trial—is discussed against the backdrop of the body of literature exploring the role of probability in theoretical and doctrinal accounts of evidence and proof, on one hand, and the role of statistical evidence in the litigation process, on the other. Understanding the nature of super-risk contributes insight to both aspects of this literature. With regard to theoretical and doctrinal accounts of the evidentiary proof process, super-risk provides a conceptual challenge for conceiving of legal proof or evidence doctrine (e.g., standards of proof) in probabilistic terms.¹⁰ With regard to statistical evidence at trial (and other litigation stages), super-risk illustrates a particular inferential problem facing legal decision-makers. This inferential problem may lead not only to erroneous decisions in individual cases, but also to systematic errors in cases of “aggregate” litigation, such as class actions, multidistrict litigation, and proposals for “trial sampling” (most notably in the context of mass torts).¹¹

Part I of the Article discusses Knight’s distinction between risk and uncertainty. Part II illustrates how risk is measured in financial decision-making. Part III illustrates how similar methods are applied to legal adjudication. Part IV details the general features of super-risk, and Part V explains how and why super-risk arises.

I. KNIGHT’S DISTINCTION

Economists and social scientists oftentimes distinguish between “risk” and “uncertainty” in decision-making.¹² The distinction is profoundly important, but the terminology is confusing. Generally, risk and uncertainty are related concepts, not distinct ones—what makes something risky is that there is uncertainty about how it will turn out. This Article will refer to “decision-making under risk” as “actuarial decision-making,” for reasons that will emerge in the discussion below.

Two examples are offered to illustrate Knight’s distinction between the two related concepts. In the first example, a woman is agonizing over whether to leave her job and start a business. In the second example, an insurer is calculating the premium for a life insurance policy with the insured’s spouse and children as the beneficiaries. The decisions of the woman and the insurer illustrate different forms of risk. For instance, the insurer can, in principle, quantify the probability of

10. In this respect, super-risk joins a list of other conceptual problems with probabilistic accounts. See Ronald J. Allen, *Factual Ambiguity and a Theory of Evidence*, 88 Nw. U. L. REV. 604 (1994); ALEX STEIN, *FOUNDATIONS OF EVIDENCE LAW* (2005); LARRY LAUDAN, *TRUTH, ERROR, AND CRIMINAL LAW: AN ESSAY IN LEGAL EPISTEMOLOGY* (2006).

11. See Richard A. Nagareda, *Class Certification in the Age of Aggregate Proof*, 84 N.Y.U. L. REV. 97 (2009); Alexandra D. Lahav, *The Case for “Trial by Formula,”* 90 TEX. L. REV. 571 (2012); Edward K. Cheng, *When 10 Trials are Better than 1000: An Evidentiary Perspective on Trial Sampling*, 160 U. PA. L. REV. 955 (2012).

12. KNIGHT, *supra* note 6.

different potential outcomes, and is thus able to predict what will occur.¹³ Pursuant to Knight's distinction, the insurer's ability to assign numerical probabilities to the different outcomes represents a case of decision-making under risk.¹⁴ Life insurers have access to actuarial tables with, for example, huge databases of mortality rates.¹⁵ Because of these tables, an insurer writing a life insurance policy is in a much stronger position (information-wise) than the person buying the policy. With an actuarial table in hand, an insurer can calculate a premium that would on average be profitable for her company, given the likelihood that people like you will die within the term of the policy and the payout that would then have to be made.¹⁶

Decision-making under "uncertainty" includes cases where decision-makers are not in a position to quantify the probabilities of different outcomes.¹⁷ Decision-makers are not necessarily clueless about the likelihood of potential outcomes (although they might be), but they do not know how certain any of the possible outcomes are. The first example illustrates decision-making under uncertainty: staying in her job may be a predictable outcome, but several factors must be considered in order to quantify the likelihood of this happening. For instance, the likely financial health of the company, the possibility of a corporate takeover, and whether she can find another job must be determined before the probability of any one outcome can be calculated. In short, in decision-making under uncertainty, such as the woman's dilemma, there are too many unknown variables to consider in order to accurately predict the likelihood of potential outcomes.¹⁸

13. See JOHN MAYNARD KEYNES, *A TREATISE ON PROBABILITY* 71–77 (1921) (referring to the amount of evidence or information on which such probabilities are based as their "weight"). Both weighty and non-weighty probabilities can give rise to super-risk. More generally, super-risk can arise whether or not a decision is epistemically justified.

14. KNIGHT, *supra* note 6, at 19–20.

15. STEPHEN SENN, *DICING WITH DEATH: CHANCE, RISK, AND HEALTH* (2003).

16. This type of decision-making is not limited to the insurance context; rather it is a staple of decision-making throughout the sciences, medicine, finance, and public policy. See Dawes et al., *supra* note 4, at 1668 (arguing in favor of this mode of decision-making generally, and explaining that it can be employed in any domain for which the relevant data are available: "[a]ctuarial output statements, or conclusions, can address virtually any type of diagnosis, description, or prediction of human interest"; to be "truly actuarial," judgments "must be both automatic . . . and based on empirically established relations").

17. KNIGHT, *supra* note 6.

18. Some statisticians and decision theorists, particularly those who follow strict Bayesianism, believe that a rational agent should always be able to assign numerical probabilities to outcomes. See LEONARD J. SAVAGE, *THE FOUNDATIONS OF STATISTICS* (1954) (the *locus classicus* of strict Bayesianism). Even as an unattainable ideal, however, it does not seem desirable to require rational agents to have unique probability functions. This is one lesson to be derived from the Ellsberg paradox. See Daniel Ellsberg, *Risk, Ambiguity, and the Savage Axioms*, 75 Q. J. ECON. 643 (1961). See also Gärdenfors & Sahlin, *supra* note 8; Isaac Levi, *On Indeterminate Probabilities*, 71 J. PHIL. 391 (1974); Isaac Levi, *The Paradoxes of Allais and Ellsberg*, 2 ECON. & PHIL. 23 (1986) [hereinafter Levi, *Paradoxes*]; Isaac Levi, *Compromising Bayesianism: A Plea for Indeterminacy*, 25 J. STAT. PLANNING & INFERENCE 347 (1990).

II. MEASURING RISK IN FINANCIAL MARKETS

Now consider a third example: an investor thinking about buying shares in a particular company. How should we think about this type of decision: Is it similar to the insurer or the woman deciding whether to quit her job? The orthodoxy in the financial world is that investment decisions are actuarial. There are many different models for valuing companies and portfolios. These models have built-in tools for measuring the different types of risk that an investor faces.

Many of these models are variations on something called the “Capital Asset Pricing Model” (“CAPM”).¹⁹ The CAPM is a way of calculating how much investors should be compensated for the risk they run in buying a particular financial asset. The basic idea is very intuitive. The more risk you run, the greater the return that you are justified in expecting from the asset. And the return that you can expect from the asset needs to be reflected in its price.

In order to apply this intuitive idea, we need a way of measuring risk. In the CAPM, risk and return are intimately connected. A correctly priced asset will yield a return that reflects the risk the investor is taking on. The greater the risk, the higher the return. The lower the risk, the lower the return. The first step in understanding the CAPM is to understand two fundamental benchmarks. These benchmarks are the starting points for working out how much you need to be compensated for holding a particular asset. They are:

Benchmark 1 = the risk-free rate of return. This is the return that an investor can expect to receive without running any risk at all. The risk-free rate of return is standardly measured by interest rates on US Treasury bills of appropriate duration, or some equally secure investment.

Benchmark 2 = the expected market rate of return. This is the return that the market as a whole is expected to yield. This might be measured in terms of the expected return of some asset designed to track the behavior of the entire stock market—an index fund tracking the Standard and Poor’s 500 in the United States, for example, or the Financial Times All Share Index in the United Kingdom.

If we subtract Benchmark 1 from Benchmark 2, then we get what is known as the *market risk premium*. This is the extra return that we can expect for taking on the risk associated with holding the market, as opposed to sitting safely with Treasury bills, or government-guaranteed savings accounts. If, for example, the return on a two-year Treasury bill is 3% and the expected return on, say, the Barclay’s corporate bond

19. The Capital Asset Pricing Model was developed independently in William F. Sharpe, *Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk*, 19 J. FIN. 425 (1964); John Lintner, *The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets*, 47 REV. ECON. & STAT. 13 (1965); and Jan Mossin, *Equilibrium in a Capital Asset Market*, 34 ECONOMETRICA 768 (1966). See also PETER L. BERNSTEIN, *CAPITAL IDEAS EVOLVING* (2007) (provides an accessible overview of the CAPM).

index is 7%, then the market risk premium is 4%. In essence, valuing a particular risky asset is identifying the correct risk premium for that asset. We know what the risk-free rate of return is—we know what return we can get without running any risk at all. What we need to know is how much extra return to expect from holding the risky asset.

One obvious question to ask is: How risky is this asset relative to the market as a whole? If the asset is less risky than the market, then the risk premium for the asset should be less than the market risk premium. If, on the other hand, the asset is riskier than the market, then it should have a higher risk premium.

This is where the elegance of the CAPM comes in. It gives a way of measuring the riskiness of a given asset relative to the riskiness of the market as a whole. Once we do that we can use the market risk premium to work out what return we should expect from our risky asset. The key idea behind measuring risk in the CAPM is that the risk of an asset is measured in terms of its volatility. The more volatile an asset, the riskier it is. Imagine two stocks, both of which increase in price by 5% over the course of a year. Suppose that the price of one stock increases steadily over the year, while the other gyrates wildly over the same period, going from a high of +15% to a low of -15%. It is natural to think that the second stock is riskier than the first. For one thing, the more volatile it is the greater the chance that if you need to cash in early you will end up taking a loss.

If we know how volatile the market as a whole is, and also how volatile the risky asset is, then we can work out the volatility of the asset relative to the volatility of the market. If the asset is more volatile than the market, then it is riskier than the market—and if less volatile, then less risky. This is the asset's *relative volatility*.

A measure of relative volatility is almost enough to allow us to work out the return we should expect from our asset—that is, the percentage of the amount invested that an investor holding that asset can expect to gain or lose. Intuitively, the expected return should be determined by the market risk premium and the volatility. If the asset is more volatile than the market, then the expected return should be higher than the market risk premium—and if it is less volatile, then it should be lower. In other words, we should be paid more for holding an asset that is riskier than the market than we get paid for holding the market.

The only thing missing is a way of representing how closely related the risky asset is with the market as a whole. This is sometimes called the *correlation coefficient* between the return on the asset and the return on the market. The closer the correlation, the more importance we need to attach to the relative volatility. The correlation can be represented by a number between 1 and -1. If the correlation is 1, then the asset and the market are perfectly correlated: they move up and down together. If the correlation is -1, then they are inversely correlated: one goes up exactly when the other goes down.

Combining relative volatility and the correlation coefficient results in what is known as the *beta* of a risky asset. The beta of a stock is

standardly taken to be a measure of how risky the stock is relative to the market as a whole. This allows us to calculate the return we should expect from the asset we are trying to value. In order to calculate the expected return, we start with the risk-free rate—since that's what we can get anyway. We then add to it the rate that we get by multiplying the market risk premium by the asset's beta.

That, in essence, is the thinking behind the Capital Asset Pricing Model. The basic ideas are:

1. The price of an asset is fixed by the return that can be expected from it.
2. The expected return is fixed by how risky it is relative to the market as a whole.
3. The asset's riskiness relative to the market as a whole is fixed by
 - a) How volatile it is relative to the volatility of the market.
 - b) How correlated its returns are with the returns of the market.

We have presented the CAPM as informally as possible. But the reason the model is so powerful and has been so influential is that there are ways of assigning numbers to each of the crucial parameters—to the relative volatility of the risky asset and to its correlation coefficient with the market.

How do we calculate relative volatility and correlation coefficients? We only have access to historical data, so that's all we can use. The relative volatility of two assets is typically calculated by comparing their individual volatility ratings over a given historical period. One very common way of looking at how volatile an asset is over a particular period of time is to look at how much it diverges from its average price during that period. This is what statisticians call *standard deviation*.

Suppose, for example, that you are interested in the volatility of shares of Main Street Bank over a twelve-month period. The first thing to do is work out the average price of a share for that period. Then you need to decide how fine-grained you want your volatility analysis to be. Are you interested in monthly fluctuations? Weekly? Daily? Or perhaps even more fine-grained than that—hourly, for example?

Suppose that you are interested in day-to-day volatility. Then for each day in the year you can measure the gap between the price on that day (perhaps the price at closing) and the average price over the twelve-month period. If on most days the share price is only a little above or below the average price, then Main Street Bank stock is not very volatile. But if there are many days when the share price is a long way above or below the twelve-month average, then it is highly volatile. Calculating the standard deviation is an easy way of measuring the degree of volatility (particularly if you have a spreadsheet program).

Suppose you know the historical volatility of Main Street Bank and the historical volatility of the market (for example, the Standard and Poor's 500) and you want to work out their relative volatility. All you need to do is to divide the first by the second. The relative volatility of

Main Street Bank and the S & P 500 over a given period is given by the ratio of their individual volatilities over that period. And so, if the volatility of a financial asset is measured by that asset's standard deviation, then we have an easy and precise way of measuring relative volatility.

The correlation coefficient is also not hard to calculate, although the calculations are a little harder to describe. As mentioned earlier, the result of the calculations is a number between 1 and -1. If the correlation coefficient is 1 then the two assets are perfectly correlated. Perfectly correlated assets move perfectly in step with each other. When one goes up or down, so does the other. If the correlation coefficient is -1, in contrast, they are exactly inversely correlated. This means that when one goes down by a certain amount, the other goes up by that same amount.

So, putting all this together, we can see how the CAPM makes certain types of investment decisions look very much like instances of actuarial decision-making. We can assign numbers to measure the risk and use those numbers to arrive at a view of the correct price for the asset. And the CAPM is just an example. There are many other models and variations upon models, all with the same aim—to find a way of measuring risk and to show how those measurements can be used to make investment decisions.²⁰

III. MEASURING RISK IN ADJUDICATION

Now consider a fourth example: a judge or jury finding facts at trial. How should we think about this type of decision: Is it similar to the insurer or to the woman deciding whether to quit her job? The primary risk in adjudication is an erroneous outcome. Errors in this context include both false negatives (failing to find a fact proven when it is true) and false positives (finding a fact proven when it is false). Both types of errors impose considerable costs on parties and society. The conventional theoretical account of adjudicative risk relies on the same basic conceptual tools (e.g., probability and utility) as the financial context, and it conceives of decision-making as actuarial.

The law manages risk in adjudication with burdens and standards of proof.²¹ For any disputed factual issue that is material to the determination of a lawsuit, the law assigns one party with the burden of proof and a standard by which the issue must be proven. For example, in civil cases, the plaintiff must typically prove the elements of her claim by a "preponderance of the evidence," and, in criminal cases, the prose-

20. Critical discussions of the CAPM and its predictive accuracy include Eugene F. Fama & Kenneth R. French, *The Capital Asset Pricing Model: Theory and Evidence*, 18 J. ECON. PERSP. 25 (2004); Fischer Black et al., *The Capital Asset Pricing Model: Some Empirical Tests*, in *STUDIES IN THE THEORY OF CAPITAL MARKETS* 79 (Michael C. Jensen ed., 1972).

21. Additional evidentiary rules regulate the admissibility of evidence; these rules also function, in part, to minimize and allocate the risk of error.

cution must prove the elements of each crime “beyond a reasonable doubt” (“BARD”).²²

The conventional account of how burdens and standards of proof are thought to measure and manage risk can be illustrated with probability theory. The “preponderance” standard is thought to minimize the risk of error and to allocate the risk of error roughly equally between the parties.²³ To illustrate this, we can conceive of the standards as requiring a probabilistic threshold of certainty in the truth of the proven fact, with 1 equaling certain truth and 0 equaling certain falsity. For the preponderance standard, this threshold is conventionally taken to be 0.5.²⁴ If the plaintiff proves the issue beyond 0.5, then the plaintiff wins on the issue (and the fact is considered “proven” for purposes of a legal judgment); if not, the plaintiff loses (and the fact is considered “not proven”). The risk of error is considered to be shared roughly equally because each side is sharing a roughly equal probability space.²⁵ If the fact is proven between 0 and 0.5, the plaintiff will lose even though the fact alleged may be true (a false negative). If the fact is proven between 0.51 and 1, then the plaintiff will win even though the fact alleged may be false (a false positive).

Under this account, total risk is thought to be minimized for the following reasons. Suppose a plaintiff proves an issue to 0.8 and thus wins. In 100 similar cases in which plaintiffs win, we would expect 80 of those cases to be correctly determined and 20 to be errors (similarly for 100 cases that defendants win when the issue is proven to 0.2). When the issue is proven to 0.6, we would expect 60 correct decisions and 40 errors, and so on. A “0.5 rule” is thought to minimize “total expected errors” because it specifies outcomes based on which outcome is more probable.²⁶ When the rule is above or below 0.5, additional expected errors follow. For example, with a “0.4 rule,” plaintiffs will win in cases proven between 0.4 and 0.5, even though the disputed facts are more likely to be false than true.²⁷ In 100 cases proven to 0.45—and thus all

22. On the preponderance standard, see *Grogan v. Garner*, 498 U.S. 279, 286 (1991); *Herman & MacLean v. Huddleston*, 459 U.S. 375, 390 (1983). On the BARD standard, see *In re Winship*, 397 U.S. 358 (1970).

23. *Grogan*, 498 U.S. at 286 (“[T]he preponderance-of-the-evidence standard results in a roughly equal allocation of the risk of error between litigants . . .”).

24. See, e.g., *Brown v. Bowen*, 847 F.2d 342, 345–46 (7th Cir. 1988).

All burdens of persuasion deal with probabilities. The preponderance standard is a more-likely-than-not rule, under which the trier of fact rules for the plaintiff if it thinks the chance greater than 0.5 that the plaintiff is in the right. The reasonable doubt standard is much higher, perhaps 0.9 or better. The clear-and-convincing standard is somewhere in between.

Id.

25. “Roughly” because issues right at the threshold of 0.5 result in victories for the defendant.

26. See STEIN, *supra* note 10, at 144; David Kaye, *The Limits of the Preponderance of the Evidence Standard: Justifiably Naked Statistical Evidence and Multiple Causation*, 7 AM. B. FOUND. RES. J. 487 (1982).

27. The disputed facts are considered more likely to be false under this account because the judgments are taken to conform to the axiom that the probability of a proposition and its negation equal 1. Therefore, if the probability of a disputed fact being true is, say, 0.45, then the probability that it is false will be $(1 - 0.45) = 0.55$.

plaintiff victories under a "0.4 rule"—this would result in expected results of 45 correct decisions and 55 errors (false positives). Similar conclusions would apply to a "0.6 rule": in 100 cases proven to 0.55, these cases would all result in victories for the defendants and there would be 55 expected errors (false negatives).

This analytical approach to measuring and managing the risk of adjudicative error works by making two important (and simplifying) assumptions: The assumptions are that (1) the base rate for deserving parties on either side will be equal, and (2) the quantified value of the evidence matches reality. If either of these two assumptions does not hold, then the actual errors produced may diverge significantly from the expected results. For example, if all civil plaintiffs who go to trial have true claims and deserve to win, then the only actual errors that can occur are false negatives (no matter where the burden and standard of proof is set, and vice versa if none deserve to win). Moreover, if cases involve poor evidence or poorly understood evidence, then additional actual errors may follow in either direction. For example, if no plaintiff can acquire sufficient evidence to surpass a 0.5 threshold, then all plaintiffs will lose, regardless of how many deserve to win on the merits; likewise, all plaintiffs may win if it is easy to generate evidence that surpasses this threshold, regardless of the merits of the case.

The discussion thus far has focused on one of the twin aspects of a decision-theoretic approach to adjudicative risk: probability. The other primary component—utility—also factors into the analysis. The analysis above regarding the preponderance standard assumed that the cost associated with each type of error is roughly equal (and likewise for the benefits associated with each type of correct judgment). If these costs are not equal, then a higher or lower probabilistic threshold will be necessary in order to minimize the expected error costs.²⁸ One approach along these lines would be to calculate an optimal standard on a case-by-case basis based on the utilities associated with each of the four possible trial outcomes (false positive, false negative, true positive, and true negative).²⁹ Although the law considers the relative costs of errors in setting the standard, it generally rejects such a fine-grained approach,³⁰ adopting instead a general taxonomy based on the types of cases: the high "beyond a reasonable doubt" standard for the elements

28. The BARD standard in criminal cases is justified on the ground that false positives (convicting the innocent) are more costly than false negatives (acquitting the guilty). See *In re Winship*, 397 U.S. 358, 371 (1970) (Harlan, J., concurring); *Addington v. Texas*, 441 U.S. 418, 423–24 (1979) ("In the administration of criminal justice, our society imposes almost the entire risk of error upon itself. This is accomplished by requiring under the Due Process Clause that the state prove the guilt of an accused beyond a reasonable doubt."); John Kaplan, *Decision Theory and the Factfinding Process*, 20 STAN. L. REV. 1065 (1968).

29. See Larry Laudan & Harry D. Saunders, *Re-Thinking the Criminal Standard of Proof: Seeking Consensus about the Utilities of Trial Outcomes*, 7 INT'L COMM. EVID. 3–4 (2009).

30. *Santosky v. Kramer*, 455 U.S. 745, 757 (1982) (quoting *Mathews v. Eldridge*, 42 U.S. 319 (1976)) ("Standards of proof, like other 'procedural due process rules[,] are shaped by the risk of error inherent in the truth-finding process as applied to the generality of cases[]"') (alteration in original) (emphasis omitted).

of crimes³¹; the preponderance standard for most civil causes of action and many affirmative defenses (criminal and civil);³² and in some jurisdictions an intermediate “clear and convincing evidence” standard for civil cases in which the relative costs of errors are thought to be sufficiently unequal.³³

Higher standards of proof, such as beyond a reasonable doubt, are thus thought to skew the risk of error in one direction. If the BARD standard is conceived as a probabilistic threshold beyond, say, 0.9, then the prosecution bears the risk of error for issues proven between 0 and 0.9, which would result in acquittals. And defendants bear the risk of error for issues proven between 0.91 and 1, which would result in convictions. The prosecution thus bears a risk of error for issues proven beyond 0.5 and up to 0.9—these issues are considered “not proven” even though they are more likely to be true than false.

To illustrate this account of adjudicative risk, consider the following examples. Suppose legal fact-finders (judges or juries) must make four factual determinations: (1) whether a bus involved in a traffic accident was owned by a particular civil defendant; (2) whether a carpet fiber found at a crime scene came from a particular criminal defendant’s home; (3) the amount of drugs smuggled by a criminal defendant; and (4) the amount of injuries suffered by civil plaintiffs all of whom were injured by the same product. One primary difference between these examples and those discussed above is that adjudication typically involves inferences about what happened in the past (“postdictions”) as opposed to predictions about future events.³⁴ Despite this difference, however, the relevant inferential issues are the same.³⁵ We discuss each example in turn.

Suppose that in the bus-accident scenario the plaintiff could prove that a bus caused the accident but could not prove the color of the bus.³⁶ The defendant owns the Blue Bus Company, which operates eighty percent of the buses in the town where the accident occurred. Assuming no other evidence, conventional expected-error analysis suggests that it is therefore 0.8 likely that the defendant’s bus caused the accident. If the standard of proof is conceived as proof beyond 0.5, then the plaintiff should win. Moreover, in 100 similar cases, we would expect 80 correct decisions and 20 errors (false positives).

31. *In re Winship*, 397 U.S. 358 (1970).

32. *Grogan v. Garner*, 498 U.S. 279, 286 (1991).

33. *Addington*, 441 U.S. at 423–24.

34. Some adjudicative decisions do, however, involve predictions about future conditions (e.g., predicting future dangerousness), as well as inferences about current conditions (e.g., injuries).

35. We note that there are important differences in how people *perceive* prediction and postdiction and that these differences are relevant to some legal issues. See Ehud Guttel & Alon Harel, *Uncertainty Revisited: Legal Prediction and Legal Postdiction*, 107 MICH. L. REV. 467 (2008).

36. See Amos Tversky & Daniel Kahneman, *Evidential Impact of Base Rates*, in JUDGMENT UNDER UNCERTAINTY: HEURISTICS AND BIASES 153, 156–58 (Daniel Kahneman et al. eds., 1982).

In the carpet fiber scenario, suppose a robbery has occurred and the police locate a carpet fiber at the crime scene that did not come from any carpets in the house that was robbed.³⁷ At trial, an expert testifies for the prosecution that "based on manufacturing records, the frequency of the recovered fiber in carpets is less than 1 in 500." The expert testifies further that the fiber recovered at the scene matches a sample of carpet taken from the defendant's home. Conventional expected-error analysis determines the "probative value" of the evidence based on a concept known as the *likelihood ratio*—which measures the probability of this evidence (a match) given a hypothesis (the defendant committed the crime), divided by the probability of this evidence (a match) given the negation of the hypothesis (the defendant did not commit the crime).³⁸ In the carpet fiber example, the probability of the evidence, assuming guilt, is (we will assume) 1. The probability of the evidence, assuming innocence, is derived from the 1 in 500 number, or 0.002 (1/500), which results in a likelihood ratio of 500 to 1. Assuming this is the only evidence offered by the prosecution, the defendant's probability of guilt is 0.998 (500/501), which would overcome a decision standard of beyond 0.9. Moreover, in 1,000 similar cases, it is expected that there would be 998 correct decisions (true convictions) and 2 errors (false convictions).

These bus-accident and carpet fiber examples illustrate the conventional account of risk in adjudication. This account of risk specifies a proof threshold, quantifies the value of evidence, and determines whether the evidence surpasses the threshold. Similar risks arise when the factual finding itself requires quantification. The next two examples illustrate this phenomenon.³⁹

Suppose a criminal defendant is arrested and convicted of illegally importing and possessing heroin.⁴⁰ His sentence will depend, in part, on the total quantity of heroin that he smuggled. Suppose further that the defendant was arrested on a flight from Nigeria with 427 grams of heroin in his possession. The prosecution presents evidence that the defendant took seven other smuggling trips to Nigeria, and presents data detailing the quantity of heroin recovered from 177 other Nigerian heroin smugglers arrested at the same airport during the period of the defendant's trips. Pursuant to the sentencing guidelines, the sentencing judge must make a factual finding regarding the total amount of heroin the defendant smuggled. Based on the data and the number of defendant's trips, the judge could conclude that he smuggled

37. This example is based on the one presented in Michael O. Finkelstein & Bruce Levin, *On the Probative Value of Evidence from a Screening Search*, 43 JURIMETRICS J. 265, 266–69 (2003).

38. See Richard O. Lempert, *Modeling Relevance*, 75 MICH. L. REV. 1021, 1023–25 (1977).

39. These types of issues interact in cases of quantification. For example, a plaintiff may have to prove by a preponderance of the evidence that she suffered a quantified amount of damages. We put aside this additional complexity for purposes of the analysis to follow, as nothing turns on it.

40. See *United States v. Shonubi*, 962 F. Supp. 370 (E.D.N.Y. 1997); *United States v. Shonubi*, 103 F.3d 1085 (2d Cir. 1997).

between 1,000 and 3,000 grams of heroin, and accordingly sentence the defendant.

Suppose 3,000 plaintiffs bring lawsuits claiming that they were injured by a product negligently manufactured by the defendant.⁴¹ After the defendant is found negligent, individual damage awards are determined by “sampling” the cases, rather than holding 3,000 separate trials. Sampling requires the court to first divide the plaintiffs into different classes based on similar characteristics, and then to hold trials for a select number of “representative cases” in each class to determine the compensatory damages for plaintiffs within each class. For instance, in *Cimino v. Raymark Indus., Inc.*, the court divided the plaintiffs into five different groups based on the types of injuries alleged, and assigned all plaintiffs to one of the groups.⁴² The court then held trials for 20 “representative cases” in each group, and awarded damages for each category based on the average award in the representative cases.

IV. FROM RISK TO SUPER-RISK

The actuarial models explored in the previous Parts are based on a single basic assumption—the *projectability assumption*. According to this assumption, data about some set of individuals or events are thought to apply to another set of individuals or events. The inferences based on this assumption project into the future, the past, or the present. For instance, regarding financial decision-making, the projectability assumption posits that historical measures of key factors, such as volatility, are good indicators of future measures of those factors. In other words, the projectability assumption suggests that the historical measures will project into the future. Similarly, regarding fact-finding at trial, the projectability assumption holds that data about one set is a good guide to what happened in the past; in other words, that the evidence projects into the past. The projectability assumption also maintains that data about current conditions is a good guide to other current conditions. For example, data about injuries in “sample” cases can be used to predict injuries in other cases.

If the projectability assumption is valid, then using a model based on known variables to make decisions is an instance of actuarial decision-making. However, what is the consequence if the projectability assumption is flawed, if historical measures of volatility and correlation do not project into the future, if evidence does not project into the past, or if “sample” cases do not project to other cases? If the projectability assumption fails, then all bets are off, because the risk

41. The trial court in *Cimino v. Raymark Indus., Inc.*, 751 F. Supp. 649 (E.D. Tex. 1990), proposed such a sampling plan for asbestos litigation. *Id.* at 651–53. However, the sampling plan was reversed on appeal. See *Cimino v. Raymark Indus., Inc.*, 151 F.3d 297, 335 (5th Cir. 1998). For recent defenses of the use of such sampling plans, see Lahav, *supra* note 11 and Cheng, *supra* note 11.

42. See, e.g., *Cimino v. Raymark*, 751 F. Supp. at 653 (where the trial court divided the plaintiffs into five different classes based on the diseases they contracted from their exposure to asbestos—mesothelioma, lung cancer, other cancer, asbestosis, and pleural diseases).

measures do not apply. With regard to financial decision-making, risk measures may indicate how assets have behaved in the past, but they give no guide as to how they might behave in the future. In adjudication, risk measures would also be poor guides as to the likelihood of particular facts being true. If the projectability assumption does not hold, then investments and adjudicative decisions are similar to what Knight called "decision-making under uncertainty."

When the projectability assumption fails, however, decision-makers relying on the models are not exactly in the same position as someone who is *knowingly* making a decision under uncertainty. Someone knowingly making a decision under uncertainty knows that she cannot assign numbers to the different outcomes. The decision-maker also knows that she must take a different approach. For instance, when making financial or policy decisions, a different approach would include adopting a "minimax" strategy and choosing the option that has the least-bad worst-case scenario.⁴³ With legal adjudication, decision-makers may adopt different inferential strategies, such as comparing competing explanations of the evidence and the events rather than trying to quantify the value of the evidence and compare it with a probabilistic threshold.⁴⁴

Unfortunately—and most critically—decision-makers are often not in a position to know *at the time of decision* whether the projectability assumption holds. Someone applying the CAPM to financial decisions needs to know whether the historical measures are going to project into the future. Decision-makers can certainly measure how well the model has projected into the future up to the point of decision—but not whether the historical projectability will continue. Similarly, the legal decision-maker needs to know whether evidence will project into the past or whether "sample" cases project to other cases. Like financial decision-makers who cannot predict the future, legal adjudicators often are not in a position to "postdict" the past or determine how well the samples project to current cases. Indeed, adjudication typically presents a worse epistemic situation than financial decisions because it will eventually be known whether the historical data projected in financial decisions, whereas with adjudication, it may never be determined whether the projectability assumption failed.⁴⁵

In both contexts, decision-makers do not know whether the assumption holds at the time of decision. Therefore, they do not know

43. Daniel A. Farber, *Uncertainty*, 99 GEO. L. J. 901, 927–34 (2011) (discussing different ways to model uncertainty). Much discussion of modeling uncertainty is framed by the discussion of LUCE & RAIFFA, *supra* note 5. Other authors have examined cases that neither fit the standard model of decision-making under risk, nor count as decision-making under complete uncertainty. See, e.g., Henry E. Kyburg, Jr., *Rational Belief*, 6 BEHAVIORAL & BRAIN SCI. 231 (1983); Levi, *Paradoxes*, *supra* note 18.

44. See Michael S. Pardo & Ronald J. Allen, *Judicial Proof and the Best Explanation*, 27 L. & PHIL. 223 (2008).

45. Although technological advancements, particularly developments in DNA technology, have provided some information about errors made at trials, especially the rate of false criminal convictions, there is a lack of information regarding the accuracy of civil and criminal trials.

which kind of decision-making they are engaging in—risky or uncertain—which creates the possibility of decision-making under “super-risk.” Super-risk situations arise when there is an appearance of decision-making under risk (i.e., actuarial decision-making) and when decision-makers try to project the measures of risk, but where the measures may not in fact project. Super-risk is a special type of “epistemic risk,” where the risk is that one’s model of risk assessment is faulty.⁴⁶ More specifically, epistemic risk is not simply the risk of being mistaken or unjustified in one’s probability assignments; rather, it is the risk of using a fundamentally mistaken model of reasoning. The risk of engaging in a fundamentally mistaken kind of decision-making may indeed include the risk of relying on a mistaken probability, but the latter is not an instance of the former.

The projection of risk measures creates a pattern. If a particular financial entity, for example, falls under a pattern, then the generalization describing the pattern can be used to predict the entity’s behavior. If this occurs, and if the historical generalization is sufficiently precise, then it is a situation of actuarial decision-making. But if, on the other hand, the financial entity does not fall under the pattern, then decision-makers are in a completely different situation—a situation of decision-making under uncertainty—which is much more dangerous because it carries with it the illusion of actuarial decision-making.

Similar considerations apply to adjudication. If a particular event, party, or disputed factual issue falls under a pattern, then the generalization describing the pattern can be used to draw inferences about the particular event, party, or legal issue. These inferences may involve postdictions, predictions, or judgments about current conditions. But if, on the other hand, the event, party, or disputed issue does not fall under the pattern, then decision-makers are engaging in decision-making under uncertainty, with the added danger that it appears to be a case of actuarial decision-making. In both the financial and the legal examples, the main point is that, at the moment of decision-making, the decision-maker is acting in conditions of risk or uncertainty, but the decision-maker does not know which one, which is the ultimate danger of super-risk.

V. UNDERSTANDING SUPER-RISK

Super-risk arises from general features of certain types of decision-making that are prominent in both financial and legal decision-making situations. These features, however, potentially exist in actuarial decision-making in the contexts of the physical and social sciences, medicine, and other domains. Comparing the legal and financial contexts with aspects of these other domains further reveals the nature of super-risk and the conditions that give rise to it.

To appreciate the general features of super-risk, two concepts must be examined in the abstract: generalizations and patterns. For our pur-

46. See Gärdenfors & Sahlin, *supra* note 8.

poses, *generalizations* are descriptions of patterns, and *patterns* are regularities in the world. In financial decision-making, correlation coefficients and measures of relative volatility are generalizations describing patterns in the world. Similarly, with adjudication, statistical evidence, likelihood ratios, and sample cases are generalizations describing patterns in the world.

Projection works by placing new cases or examples under the description and inferring that the pattern will continue with the new cases or examples. This is true when we use a generalization to predict economic activity, to convict a criminal defendant, or to diagnose a medical condition.⁴⁷ When we use a generalization to predict how some object will behave, or how some event will turn out, we are assuming that the object or event is relevantly similar to the objects or events in the pattern that the generalization describes. We are also assuming this when we use a generalization to postdict behavior, or how some event turned out in the past, and when we use a generalization to infer conclusions about current conditions. These are inferences that we make all the time. It is no exaggeration to say that our daily life depends upon constantly applying what we call the projectability assumption—the assumption that patterns that we have identified will project.

Super-risk arises when there is a disconnect between two different properties in the generalizations relied upon in decision-making: the *explanatory property* and the *projectability property*. Generalizations have the explanatory property when they describe a genuinely existing pattern in the world.⁴⁸ Generalizations have the projectability property when the pattern described by the generalization could be extended to new cases and could thus be relied upon by a decision-maker to draw inferences about new cases. Super-risk arises when the generalizations relied upon have the explanatory property, but lack the projectability property.⁴⁹ This creates a particular type of inferential problem

47. See Cohen, *supra* note 4, at 55–56 (“[T]he trier of fact in a criminal court is in a situation that is analogous in some respects to that of a diagnostic physician The trier of fact, like the physician, cannot base his verdict *merely* on a high mathematical probability. . . . The point is that the reliability of a generalisation, in experimental science as also elsewhere, is best graded . . . by the variety of potentially relevant circumstances that fail to falsify it.”).

48. See *id.* at 60 (“[T]he patient’s condition may be so well understood that an evaluation in terms of mathematical probability would be superfluous, because so much of the causal process at work has been established”); BISHOP & TROUT, *supra* note 4, at 47 (distinguishing between *grounded* and *ungrounded* “statistical prediction rules”; grounded statistical prediction rules are ones for which “we have a theoretical explanation for their success”).

49. The fact that super-risk arises from some generalizations with the explanatory property is the reason that the phenomenon of super-risk is distinct from whether the underlying generalization is reliable or “weighty.” See Cohen, *supra* note 4, at 55–56. The projectability property, as we construe it, is broader than what Alex Stein has recently referred to as the “extendibility presumption.” Alex Stein, *The Flawed Probabilistic Foundation of Law and Economics*, 105 Nw. U. L. REV. 199, 221 (2011). Relying on Mill, Stein’s analysis correctly points out that probabilistic generalizations may be epistemically problematic if they fail to track genuine causal patterns in the world (the “extendibility pre-

whereby decision-making has the appearance of taking place in conditions of risk, due to the explanatory property, but is actually taking place in conditions of uncertainty, because the projectability property is absent.⁵⁰

We will now spell out in general terms how the link between the explanatory property and the projectability property could be broken. Thereafter, two particular causes of super-risk will be examined: the “reference class” problem and feedback loops.

Generalizations which have the *explanatory property*, tracking existing patterns in the world, typically hold in virtue of *causal connections* (and hence will not rest simply upon observed correlation).⁵¹ For example, consider the “Dogs of the Dow” investment strategy, which instructs an investor to renew her portfolio each year by buying the 10 stocks in the Dow Jones Industrial Average with the highest dividend yield.⁵² The strategy is based on the observation that large, blue chip companies with a high dividend yield have annually outperformed their peers with lower yields. If the generalization underpinning the “Dogs of the Dow” strategy has the explanatory property, this is so because the relatively high dividend yield plays a causal role in the “Dogs” outperforming the other companies in the index. Similarly, generalizations regarding legal fact-finding may also display the explanatory property. For example, plaintiffs in “sample” tort cases may have suffered the particular injuries and damages *because* they were exposed to and injured by the same product.

The potential for super-risk arises because, even when generalizations have the explanatory property, those generalizations might still fail to project. Consider again the “Dogs of the Dow” strategy. Any two

sumption” presumes that they do). *See id.* The extendibility presumption resembles what we call the explanatory property:

An occurrence of feature *B* in numerous cases of *A* does not, by and of itself, establish that *B* might occur in a future case of *A*. Only evidence of causation can establish that this future occurrence is probable. This evidence needs to identify the causal forces bringing about the conjunctive occurrence of *A* and *B*. Identification of those forces needs to rely on a plausible causal theory demonstrating that *B*'s presence in cases of *A* is law-bound rather than accidental.

Id. Even generalizations that possess the explanatory property and track causal relationships may nevertheless fail to project, and thus give rise to super-risk.

50. There are three types of inferential problems that are *not* instances of super-risk: (1) some decisions will be errors (this is just a feature of actuarial decision-making); (2) some generalizations relied on may be based on incorrect data; and (3) the connection between the pattern and the projected cases could be purely coincidental (this is just the plain, old-fashioned risk that you have mistaken a coincidental correlation for a causal one). *See* PEARL, *supra* note 4.

51. *See* COHEN, *supra* note 4; BISHOP & TROUT, *supra* note 4, at 47 (noting that for any ungrounded statistical rule “there may well be a neat causal explanation for its success that we don’t know yet”). For further discussion of the relationship between explanation and causation, see STATHIS PSILLOS, CAUSATION AND EXPLANATION (John Shand ed., 2002); CAUSATION AND EXPLANATION (Joseph Keim Campbell et al. eds., 2007).

52. *See* MICHAEL B. O’HIGGINS & JOHN MCCARTY, BEATING THE DOW WITH BONDS (1999) (the publication in which the “Dogs of the Dow” was first proposed). The dividend yield of a share is calculated by dividing the dividend per share by the price of the share.

companies *might* be similar along one dimension (e.g., high dividend yields), but will certainly be dissimilar along many other dimensions. A strategy like the "Dogs of the Dow" may fail with a particular company because one or more of the dissimilarities between it and other blue chip companies trump the similarity—that both are blue chip companies with high dividend yields. Assume the "Dogs of the Dow" strategy is based on a genuine generalization which captures more than just a correlation. In other words, the strategy is not merely that the companies with a high relative dividend yield have outperformed the index, but that they have outperformed the index *because* they have a high dividend yield. Nonetheless, there are other contributing factors: the companies from whose performance the 'Dogs of the Dow' strategy has been extrapolated resemble each other in some respects (including high relative yield), but not in others. The general point is that (outside the realm of fundamental physics, and perhaps even there) explanatory generalizations always rely, either implicitly or explicitly, on a *ceteris paribus* clause that allows for a range of circumstances in which the generalization can fail to hold.⁵³ Thus, generalizations could fail to project while still being explanatory when one of the potential defeaters is in play—even if the relevant generalization is explanatory in the sense that it tracks a genuine pattern in the world.

Compare this with actuarial decision-making in the context of life insurance. If an actuary says that smokers of a given age are more likely to die in the next twelve months than non-smokers of the same age, then we can be reasonably confident that smoking is the key contributing factor to the difference in life expectancy. Life insurers typically do not have to ask very many more questions before fixing a price for the policy. The reason, of course, is that they (and we) have a pretty clear view of the machinery that leads from smoking to a lower life expectancy. We understand the health consequences of smoking, and how they impact mortality rates.⁵⁴

53. The nineteenth century economist Alfred Marshall understood this point very clearly as applied to economics:

The forces to be dealt with are however so numerous, that it is best to take a few at a time; and to work out a number of partial solutions as auxiliaries to our main study. Thus we begin by isolating the primary relations of supply, demand and price in regard to a particular commodity. We reduce to inaction all other forces by the phrase "other things being equal": we do not suppose that they are inert, but for the time we ignore their activity. This scientific device is a great deal older than science: it is the method by which, consciously or unconsciously, sensible men have dealt from time immemorial with every difficult problem of ordinary life.

ALFRED MARSHALL, *PRINCIPLES OF ECONOMICS* xx (8th ed. 1920) (emphasis added). See DANIEL M. HAUSMAN, *THE INEXACT AND SEPARATE SCIENCE OF ECONOMICS* (1992). For general philosophical discussions, see JERRY A. FODOR, *Special Sciences (Or: The Disunity of Science as a Working Hypothesis)*, 28 *SYNTHESE* 97 (1974); DONALD DAVIDSON, *ESSAYS ON ACTIONS AND EVENTS* (1980); John Earman et al., *Ceteris Paribus Lost*, 57 *ERKENNTNIS* 281 (2002).

54. BISHOP & TROUT, *supra* note 4; see also that "many good examples of grounded [statistical prediction rules] come from medicine." *Id.* at 48. In discussing variables that predict prostate cancer, they explain that for three of the variables "we understand very well why . . . those variables help to reliably predict the target property." *Id.*

But the same does not hold for companies and dividend yields. There is nothing corresponding to medical understanding of how smoking contributes to cancer, emphysema, cardiovascular disease, and so on. Even if we assume that there is a historically robust “Dogs of the Dow” phenomenon, our understanding of why it might occur is fundamentally different from the actuarial case. We know that relative dividend yield is a causal factor, but we have no idea what the other causal factors might be. And this is not very surprising, given the massive differences between the companies involved. This means that we cannot be confident that the robust generalization we have identified in the past will carry over to the case to which we are trying to apply it. There remains the possibility that new cases will prove to be counterexamples, not instances—even though the generalization is tracking a causal connection.

This is why decision-making in financial markets is so often decision-making under super-risk. Decision-makers are often working with generalizations that they have good reason to think have the explanatory property. These generalizations typically make it seem that decisions are being made under actuarial conditions.⁵⁵ At the same time, however, these generalizations will always have an “all other things being equal” clause built into them, either explicitly or (more often) implicitly. In almost all interesting cases, it is impossible to spell out completely what these “other things” might be, or what counts as “being equal.”

This fundamental problem applies to each of the four legal examples as well. First, even if there is a real causal generalization underlying the sample tort cases (based on the similarity that all plaintiffs were injured by the same product), there may be many dissimilarities among the class of plaintiffs.⁵⁶ These dissimilarities may be legally significant and thus make the non-sample cases counterexamples rather than extensions of the patterns identified. Second, even if there is a robust generalization about the behavior of Nigerian heroin smugglers, the defendant in any given case, or defendants in a series of cases, may not be typical Nigerian heroin smugglers.⁵⁷ Third, even if the Blue Cab Company owns eighty percent of the cabs in town, they may hire safer drivers than other companies and its cabs may be involved in only ten percent of the cab accidents in town.⁵⁸ Finally, even if the carpet fiber found at the crime scene and the defendant’s home is rare in carpets

55. There are broader psychological issues involved also. Daniel Ellsberg drew attention to the psychological phenomenon of ambiguity aversion. See Ellsberg, *supra* note 18. There is robust evidence that decision-makers regularly breach the axioms of expected utility theory when they are dealing with probabilities that are not numerically determinate, but rather fall into a range. The fact that people typically prefer to deal with numbers rather than with ranges could be one source of the phenomenon under discussion. *Id.*

56. See *Cimino v. Raymark Indus., Inc.*, 751 F. Supp. 649 (E.D. Tex. 1990), *rev'd*, 151 F.3d 297 (5th Cir. 1998). See also Lahav, *supra* note 11; Cheng, *supra* note 11.

57. See *United States v. Shonubi*, 962 F. Supp. 370 (E.D.N.Y. 1997), *vacated*, 103 F.3d 1085 (2d Cir. 1997).

58. See Tversky & Kahneman, *supra* note 36.

manufactured nationally, the fiber may be common in carpets sold at the local carpet store.⁵⁹ As with the financial example, in each of these examples there are both (1) reasons for legal decision-makers to think they have identified generalizations describing actual patterns in the world, and (2) reasons why the generalizations may not project to the new cases.

Actuarial decision-making—assigning numbers to the probabilities of different outcomes and measuring risk—is a perfectly reasonable way to proceed, to the extent that you have evidence that you are working with patterns that have the explanatory property. But, of course, this depends upon the assumption that explanatory patterns will project. Or, to put it another way, it rests upon the tacit assumption that generalizations that have the *explanatory* property also have the *projectability* property. In cases of actuarial decision-making in the insurance context, and in many scientific and medical contexts, this assumption is well-founded.⁶⁰ Actuarial generalizations project to new cases. If you have good evidence that a particular factor (for example, smoking) has had a particular effect (such as decreasing life expectancy) in the past, then that gives you license to make predictions about the future. You can use those predictions to write your insurance policies with confidence (assuming that you write a sufficient amount of them for the statistics to work in your favor).

The link between the explanatory and projectability properties typically holds in the physical sciences as well. If we have evidence that some apparent generalization in physics, for example, fails to project to new cases, then we should suspect that the explanatory property doesn't hold. There may be some complex scientific phenomena that cannot be understood in terms of actuarial generalizations. Climate change may be an example. But much of science, particularly basic science, is actuarial (in our sense). In many types of scientific generalizations, in other words, the explanatory property and the projectability property stand or fall together.⁶¹

But generally this does not hold for financial markets and legal adjudication. Certain distinctive features of each break the link between the explanatory property and the projectability property. And so we end up in decision-making under super-risk, the situation of making decisions under what we take to be actuarial conditions when we might actually be operating in a situation of uncertainty rather than risk.⁶²

We now look at two particular sources of super-risk. These sources provide reasons why the explanatory and projectability properties may come apart in legal and financial decision-making.

59. See Finkelstein & Levin, *supra* note 37.

60. See, e.g., BISHOP & TROUT, *supra* note 4; Dawes et al., *supra* note 4.

61. Cf. ALEXANDER ROSENBERG, ECONOMICS—MATHEMATICAL POLITICS OR SCIENCE OF DIMINISHING RETURNS? (1992) (distinguishing economics from the sciences based on the former's failure to achieve similar predictive successes).

62. See, e.g., HAUSMAN, *supra* note 53.

The first is the reference-class problem.⁶³ In adjudication, inferences drawn from statistical evidence about a particular event depend on placing the event within a particular reference class, the class described by the data.⁶⁴ The reference class provides the generalization describing the pattern that will then be applied to the particular event.⁶⁵ The particular event, however, is a member of a virtually infinite number of other reference classes, each of which may provide widely different generalizations describing different patterns.⁶⁶ For example, even though the Blue Cab Company owns eighty percent of the buses in the town, suppose again that it causes ten percent of the accidents. Then, even though the “80 percent” generalization describes a real pattern in the world, it will not project to cases like the bus accident. If it did project, then in ten similar cases (all of which will result in judgments against the Blue Cab Company) we would expect eight correct decisions and two errors—what we would get instead are nine errors and one correct decision.

The reference-class problem is a general one that applies to adjudication as well as to financial decision-making. Differences between members of the class described in any given generalization, on one hand, and the target cases, on the other, point to other reference classes to which the target cases are members. Any of these differences may be a reason why the generalization fails to project to the target cases.⁶⁷ The pattern of carpet fibers in the United States may be a different pattern than carpet fibers in Anytown, USA; the pattern of Nigerian heroin drug smugglers may be different from the pattern of Nigerian heroin drug smugglers over the age of 30; the pattern of people injured by a product may be different from people with high blood pressure injured by a product; and the pattern of companies with a high dividend yield may be different from the pattern of companies with a high dividend yield that also issue fraudulent financial statements. The reference-class problem reveals one way in which generalizations may describe a real pattern in the world but fail to project to new cases.

A second way the explanatory and projectability properties can come apart is when reliance on the generalizations undermines the generalizations. This process can occur with both financial and legal decision-making. We illustrate first with a financial example, again involving the “Dogs of the Dow” strategy.

Patterns in the behavior of financial assets can be self-destructive. When investors notice the patterns and start to act upon them, their very actions can cause the patterns to cease to hold. If investors realize that the “Dogs of the Dow” will outperform their companions with a lower dividend yield, then they will adopt the investment strategy of

63. Alan Hájek, *The Reference Class Problem is Your Problem Too*, 156 SYNTHESE 563 (2007); see also Ronald J. Allen & Michael S. Pardo, *The Problematic Value of Mathematical Models of Evidence*, 36 J. LEGAL STUD. 107 (2007).

64. Hájek, *supra* note 63, at 564.

65. *Id.*

66. *Id.* at 567.

67. *Id.* at 565.

buying those underappreciated and overperforming companies. This will drive their share price up and the end result will be that they stop being underappreciated and overperforming. There is some evidence that this is what actually happened to the "Dogs of the Dow" strategy.⁶⁸

This phenomenon works both ways. If you think you have spotted an anomaly in the market, then you would be wise to consider the probability of other people spotting the very same anomaly. Unless you are doing something illegal, the information you have is shared by other investors. So you should consider the possibility that there isn't an anomaly there at all because, if there were one, the market would already have closed it down. The efficient-markets hypothesis generalizes this type of argument.⁶⁹ In a sense this phenomenon is not very mysterious. It is simply an example of what is often called the principle of no arbitrage—the principle that markets will close down any opportunity for a risk-free profit as soon as it appears.⁷⁰ There is a related phenomenon in social policy. Goodhart's Law states that when policy-makers explicitly target a given economic indicator (such as a given level of inflation) this changes its indicator function (e.g., what it tells us about unemployment).⁷¹

Neither the efficient-markets hypothesis nor Goodhart's Law is by *itself* a source of super-risk, but the "Dogs of the Dow" example illustrates a fundamental fact about financial markets that does provide a source of super-risk. Financial markets move the way they do because of how investors behave. Investors behave the way they do because of how financial markets have moved in the past, and how they expect them to move in the future. This gives rise to a feedback loop that George Soros has labeled "reflexivity."⁷² Here is how Soros describes reflexivity:

In situations that have thinking participants, there is a two-way interaction between the participants' thinking and the situation in which they participate. On the one hand, participants seek to understand reality; on the other, they seek to bring about a desired outcome. The two functions work in opposite directions: in the cognitive function reality is the given; in the participating function, the participants' understanding is the constant. The two functions can interfere with each other by rendering what is supposed to be given, contingent. I call the interference between the two functions "reflexivity." I envision reflexivity as a feedback loop between the participants' understanding and the situation in

68. Grant McQueen et al., *Does the "Dow-10 Investment Strategy" Beat the Dow Statistically and Economically?*, 53 FIN. ANALYSTS J. 66 (1997).

69. Eugene F. Fama, *Efficient Capital Markets: A Review of Theory and Empirical Work*, 25 J. FIN. 383 (1970).

70. See Hal R. Varian, *The Arbitrage Principle in Financial Economics*, 1 J. ECON. PERSP. 55 (1987).

71. Charles A.E. Goodhart, *Problems of Monetary Management: The U.K. Experience*, in PAPERS IN MONETARY ECONOMICS (National Bureau of Economic Research, 1975).

72. GEORGE SOROS, *THE ALCHEMY OF FINANCE: READING THE MIND OF THE MARKET 2* (2003).

which they participate, and I contend that the concept of reflexivity is crucial to understanding situations that have thinking participants.⁷³

Soros illuminates the reflexivity of financial markets (and has also profited from it rather impressively). The reflexivity that Soros describes is a fundamental source of super-risk.

Detecting historical patterns can lead investors to behave in certain ways. This is the first part of Soros's feedback loop between information and participation. But then, and this is the second part of the loop, that behavior can fundamentally alter the basis for the generalization—participation can change the information base. The run-up to the recent financial crisis provides a very clear illustration of this. As is by now very well-documented, the pricing of mortgage-backed securities—securities that are backed by packages of mortgages on residential or commercial properties—was determined by the risk of the underlying assets defaulting. This risk was typically measured with reference to historical default rates. The models for pricing these securities based on the historical data allowed the securities to be sliced and diced into different layers or tranches, each yielding a different rate of return corresponding to a different level of perceived risk.

What happened next is well-known. Banks and mortgage companies started both selling mortgages to individuals who would never previously have qualified for loans and offering new types of mortgage products. They did so, in part, by assuming the *projectability* of historical default rates into the future.⁷⁴ This behavior ensured that the historical rates did not project. Default rates ended up far above their historical levels, and mortgage-backed securities turned into “toxic assets.”⁷⁵

This illustrates how the reflexivity of financial markets can lead generalizations and patterns to undermine themselves. Mortgage companies and securitization bankers correctly identified a robust historical pattern of repayment and default in residential and commercial mortgages. They assumed that this pattern would project into the future and fed it into their models. This had the direct consequence that the pattern undermined itself. The mortgage industry eventually discovered that they were working with a pattern that had the *explanatory* property but not the *projectability* property.

This type of reflexivity arises in legal contexts as well. The fact that a generalization is relied on in one case is a reason for parties (particularly, so-called “repeat players”) to take advantage of or to seek to undermine the generalization in future cases. Evidentiary rules and

73. *Id.*

74. In its report on the causes of the financial crisis, the U.S. Financial Crisis Inquiry Commission noted that, “[f]inancial institutions and credit rating agencies embraced mathematical models as reliable predictors of risks, replacing judgment in too many instances. Too often, risk management became risk justification.” FIN. CRISIS INQUIRY COMM’N, THE FINANCIAL CRISIS INQUIRY REPORT xix (2011). See also POSNER, *supra* note 9 (noting the significant role played by problematic mathematical models of risk in the financial crisis).

75. FIN. CRISIS INQUIRY COMM’N, *supra* note 74, at 480.

legal judgments affect primary (i.e., non-litigation) behavior and changes in primary behavior affect generalizations and their projectability in future cases.⁷⁶ If the hypothetical Blue Cab Company lost one case based on evidence of the number of buses in the town it owns, we can imagine it would take steps to undermine the projectability of this generalization to any future cases—by, for example, hiring safe drivers and generating evidence that tracks the exact locations of their buses or perhaps placing cameras on the buses. As this example illustrates, “undermining” the projectability of generalizations through such reflexivity may be a good thing in some instances. For another example, if the police can use data about patterns of criminal behavior to predict where certain crimes are likely to occur, then deploying more police officers to those areas may indeed reduce the current patterns and undermine the generalizations describing them.⁷⁷ Likewise, in the medical context, the presence of certain risk factors may be a good predictor for developing a particular disease, but possessing these factors may thereby cause people to change their behavior to decrease the odds of developing the disease. Nevertheless, this type of reflexivity provides a fundamental way in which the explanatory and projectability properties may diverge and thus is a source of super-risk.

CONCLUSION

Adjudication is a process of drawing uncertain inferences about the past, present, or future. Investment is a process of making educated guesses about the future. The stock in trade of both enterprises is, as it is throughout the physical and social sciences, perceived similarities and potential patterns. The complexity of the financial world and the ingenuity of its inhabitants have created almost endless possibilities for perceiving similarities and identifying patterns. The complexity and the endless variety of issues that give rise to adjudication create similar possibilities. The patterns identified in each domain can often be made numerically precise.

Therefore, from the perspectives of law and finance, decision-making can all too easily have the appearance of being actuarial, akin to insurance and some aspects of the physical sciences. But that is often a mistake. Instead, it is decision-making under super-risk. This occurs when decision-makers identify precise patterns—e.g., the past behavior of financial entities or statistical evidence about a class of events or individuals. These patterns can be made numerically precise using a variety of tools. For the financial analyst, these include expected return, beta, dividend yield, and so on. For the legal fact-finder, these include statistical evidence, likelihood ratios, Bayes’ Theorem, and probabilistic standards of proof. But the apparent numerical precision in these domains

76. Gideon Parchomovsky & Alex Stein, *The Distortionary Effect of Evidence on Primary Behavior*, 124 HARV. L. REV. 518 (2010); Chris William Sanchirico, *Character Evidence and the Object of Trial*, 101 COLUM. L. REV. 1227 (2001).

77. Andrew Guthrie Ferguson, *Predictive Policing and Reasonable Suspicion*, 62 EMORY L.J. 259 (2012).

may flatter to deceive. Decision-makers are often in a position of super-risk because, while they may be basing decisions on genuine explanatory patterns, they have no way of knowing whether these patterns have the projectability property.

Ordinary risk can be measured and tamed. An insurer who writes a single insurance policy may get badly burned. But if insurers write enough policies, then on balance they are likely to come out ahead. Can super-risk be tamed in this way? Our instinct is that it cannot. But that is a topic for another time. Our principal aim in this Article has been to introduce the important, but thus far neglected, concept of super-risk.