2. Behavioral probability

Alex Stein

1. INTRODUCTION

This chapter examines experimental studies that identify misalignments between ordinary people's decisions under uncertainty and the rules of mathematical probability.[1] These studies use mathematical probability as a criterion for rational decisions. Based on this criterion, the studies tag people's deviations from mathematical probability as irrational (or as boundedly rational). The studies also identify those deviations' recurrent patterns and develop a taxonomy for describing people's probabilistic mistakes. Under this taxonomy, those mistakes include "representativeness," "availability," "base-rate neglect," and suppression of the product rule.

Representativeness is a person's preference of familiar scenarios over statistical data (Kahneman 2011, pp. 146–55). Availability is an individual's overestimation of the probabilities attaching to events that fall within her experience or easily come to mind (Kahneman 2011, pp. 129–36). Base-rate neglect is a probability assessment that fails to consider general distributions of relevant events (Kahneman 2011, pp. 166–74). Suppression of the product rule is a person's failure to treat a compound event (events *A* and *B* occurring simultaneously) as less probable than each of its components (*A* or *B*) (Kahneman 2011, pp. 156–65).

Arguably, these mistakes lead to erroneous decisions that adversely affect people's welfare. Behavioral economists[2] argue that the government should step in to prevent these erroneous decisions. Specifically, they recommend the following legal reforms: mandatory supply of information to error-prone individuals,[3] soft choice-architecture,[4] and regulatory intervention that will prevent and correct people's probabilistic mistakes.[5] Areas targeted by these recommendations include accidents and risk regulation, consumer

---

[1] The most significant of those studies are reported and analyzed in Kahneman (2011). Written by the discipline's founding father, this book is sure to become a canonical text on behavioral probability.

[2] This designation includes not only economists, but also psychologists investigating the ways in which people reason and make decisions.

[3] See, e.g., Bar-Grill and Stone (2009) using behavioral economics to propose expansive disclosure requirements in connection with cellular service agreements. But see Ben-Shahar and Schneider (2011) criticizing the ongoing expansion of disclosure requirements.

[4] See Thaler and Sunstein (2008) introducing the "choice architecture" method, understood as governmental manipulation of individuals' menu of choices in a manner that nudges those individuals to take the desired action.

[5] For a summary of regulatory initiatives driven by behavioral economics and an analytical framework for regulation premised on subjects' bounded rationality, see Vandenbergh, Carrico and Schultz Bressman (2011, pp. 763–78).

agreements, business contracts, credit and lending, employment, insurance, prenuptial agreements, and adjudicative fact-finding.[6]

Studies surveyed herein have been carried out by Daniel Kahneman, Amos Tversky and other behavioral economists. These studies form a distinct field of inquiry, identified here as "behavioral probability." Behavioral probability is part of a more comprehensive area of study: behavioral economics. Behavioral economics is a discipline that encompasses behavioral probability along with experimental and empirical studies of people's assessments of utility. Behavioral economics has been immensely successful as a general discipline: it has influenced many studies of economics, finance, and law (Vandenbergh, Carrico, and Schultz Bressman 2011; Bar-Gill and Warren 2008; Eisenberg 1995; Jolls, Sunstein, and Thaler 1998; Kahan 2010; Rachlinski 1998; Sunstein 1986; Williams 2009; Zamir 1998).

This chapter is organized as follows. Section 2 examines the rules of mathematical probability that the studies surveyed herein use as a benchmark for rationality. Section 3 juxtaposes these rules against people's hardwired habit of understanding the world in terms of causes and effects. This juxtaposition identifies a serious tension between mathematical probability and people's causal understanding of the world. I show that this causal understanding is not indicative of people's irrationality (or bounded rationality). Far from irrational, people's causal understanding of the world has its own probabilistic framework, identified as inductive, or Baconian, probability. Section 4 uses these insights to revisit the experiments carried out by Kahneman, Tversky, and other behavioral economists and tendered as a proof of people's probabilistic failures. I demonstrate that these experiments do not establish that people are probabilistically irrational (or boundedly rational). In fact, I show that some of these experiments are methodologically flawed.

[6] See, e.g., Bar-Gill and Warren (2008) identifying and calling for regulatory correction of people's over-optimism in consumer credit agreements; Eisenberg (1995) identifying over-optimism in people's liquidated damage undertakings, prenuptial agreements and other areas of contract and commending a regime that authorizes courts to modify contractually prearranged payments and waivers; Jolls, Sunstein and Thaler (1998, pp. 1522–28) identifying and calling for regulatory correction of hindsight biases in courts' determinations of negligence, environmental torts, punitive damages, and non-obviousness of patented inventions; Kahan (2010, pp. 1623–25) describing the effect of hindsight bias on fact-finding in adjudication; Rachlinski (1998) identifying the presence of hindsight bias in courts' ascertainments of parties' compliance with ex-ante norms and commending legal rules that counteract this bias; Sunstein (1986, pp. 1167–68) identifying and calling for regulatory correction of base-rate neglects in people's decisions about risk and insurance, contractual undertakings, and their own employment termination prospects; Williams (2009) identifying and calling for regulatory correction of people's base-rate neglects and resulting overconfidence in marriage-related and employment agreements and in credit card borrowing; Zamir (1998, pp. 269–70) identifying and calling for regulatory correction of people's base-rate neglects and availability bias in savings and credit decisions.

2. MATHEMATICAL PROBABILITY: LANGUAGE AND EPISTEMICS[7]

The best way to understand mathematical probability is to perceive it as a language that describes the facts relevant to a person's decisions. Like all languages that people use in their daily interactions, the probability language has a set of conventional rules. These rules determine the meanings, the grammar, and the syntax of probabilistic propositions. Compliance with these rules enables one person to form meaningful propositions about probability and communicate them to other people.

The probability language differs from ordinary languages in three fundamental respects: scope, parsimony, and abstraction. First, ordinary languages have a virtually unlimited scope, as they promote multiple purposes in a wide variety of ways. People use those languages in communicating facts, thoughts, ideas, feelings, emotions, sensations, and much else. The probability language, in contrast, has a much narrower scope because it only communicates the reasoner's epistemic situation or balance of knowledge versus ignorance. The reasoner uses this language to communicate what facts she considers relevant to her decision and the extent to which those facts are probable. Second, ordinary languages have rich vocabularies.[8] The probability language, by contrast, is parsimonious by design: it uses a small set of concepts to describe multifarious events in a standardized mode. This mode establishes a common metric for all propositions about the probabilities of uncertain events. This metric creates syntactical uniformity in the probability language and makes it interpersonally transmittable. Finally, because a person usually needs to deal with more than one uncertain event, she needs a uniform set of abstract concepts by which to relate one probability estimate to another and to integrate those estimates into a comprehensive assessment of probability.

These attributes of the probability language account for its high level of abstraction, uncharacteristic of any ordinary language. To maintain the required parsimony and conceptual uniformity, the probability language uses mathematical symbols instead of words. Those symbols allow a person to formulate her assessments of probability with precision. This precision, however, is purchased at a price: the comprehensive trimming of particularities and nuances that characterize real-world facts. The scope of each assessment's meaning and applicability thus becomes opaque and at times indeterminable. This tradeoff—precise language for a weak epistemic grasp—is a core characteristic and the core problem of mathematical probability. The two components of this tradeoff stand in an inverse relationship to each other. To be able to formulate her probability assessments with precision, a person must get rid of untidy concepts, downsize her vocabulary, and abstract away the multifaceted nuances of the real world. All this weakens the person's epistemic grasp of the real world. As a result, her abstract, numerical estimates will say hardly anything informative about concrete events that unfold on the ground. To have a strong epistemic grasp of the factual world, a person has to be wordy: she must utilize a rich vocabulary and loosen her conceptual precision.

---

[7] This section is based on Stein (2011).

[8] See, e.g., *The Oxford English Dictionary* (1989) a 20-volume dictionary that explains the meanings of over 600,000 words originating from approximately 220,000 etymological roots.

2.1 The Language of Mathematical Probability

The mathematical probability system designates the numerical space between 0 and 1 (the algebraic equivalents of 0% and 100%) to accommodate every factual scenario that exists in the world:[9]

0 1

This space accommodates two propositions that are factually certain:

Proposition A: The probability that one of all the possible scenarios will materialize equals 1.

Proposition B: Correspondingly, the probability that none of all the possible scenarios will materialize equals 0.

These propositions are tautological. The first proposition essentially says, "Something will certainly happen." The second makes an equally vacuous attestation: "There is no way that nothing will happen." All other propositions occupying the probability space are meaningful because they describe concrete events that unfold in the real world. These meaningful propositions are inherently uncertain. There is no way of obtaining complete information that will verify or refute what they say. Consequently, the probability of any concrete scenario is always greater than zero and less than one. More precisely, the probability of any concrete scenario, $P(S)$, equals one minus the probability of all factual contingencies in which the scenario does not materialize: $P(S) = 1 - P(\text{not} - S)$. This formula is called the "complementation principle."[10]

To illustrate that principle, consider a random toss of a coin. The coin is unrigged: its probability of landing on heads is the same as its probability of landing on tails. Each of these probabilities thus equals 0.5. The two probabilities divide the entire probability space among themselves. The coin's probability of landing on either heads or tails equals 1, and we already know that this proposition is vacuous or tautological.

0.5

0 1

This illustration does not address the key question about the coin. What does "unrigged" mean? How does one know that this specific coin is equally likely to land on heads or on tails? This important question focuses on the epistemics of mathematical probability,

[9] My discussion simplifies Kolmogorov's classic definition of the "probability space" (Kolmogorov 1956).

[10] See Cohen (1989, pp. 17–18, 56–57) stating and explaining the complementation principle.

discussed in Section 2.2 below. My present discussion only addresses mathematical probability's syntax and semantics. For that reason, I assume for now that the two probabilities are equal. The coin's probability of landing on tails, as opposed to heads, or vice versa, is deemed to be 0.5.

We are now in a position to grasp the second canon of mathematical probability: the "multiplication principle" or the "product rule" (see Cohen [1989, pp. 18–21], stating and explaining the multiplication principle). The multiplication principle holds that the probability of a joint occurrence of two mutually independent events, S_1 and S_2, equals the probability of one event multiplied by the probability of the other. Formally: $P(S_1 \,\&\, S_2) = P(S_1) \times P(S_2)$.

My coin example makes this principle easy to understand. Consider the probability of two successive tosses of an unrigged coin landing on heads. The probability that the first toss will produce heads, $P(S_1)$, equals 0.5. The probability that the second toss will produce heads, $P(S_2)$, equals 0.5 as well. The first probability occupies half of the entire probability space, while the second—as part of the compound, or conjunctive, scenario we are interested in—occupies half of the space taken by the first probability. The diagram below shows this division of the probability space:

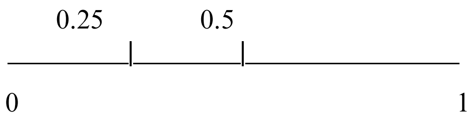

The complementation and multiplication principles are the pillars of the mathematical system of probability. All other probability rules derive from these principles. Consider the "disjunction rule" (see Kneale [1949, pp. 125–26] stating and explaining the disjunction rule) that allows a person to calculate the probability of alternative scenarios, denoted again as S_1 and S_2. This probability equals the sum of the probabilities attaching to those scenarios, minus the probability of the scenarios' joint occurrence. Formally: $P(S_1 \text{ or } S_2) = P(S_1) + P(S_2) - P(S_1 \,\&\, S_2)$. Here, the deduction of the joint-occurrence probability, $P(S_1 \,\&\, S_2)$, prevents double counting of the same probability space. The probability of each individual scenario, $P(S_1)$ and $P(S_2)$, occupies the space in which the scenario unfolds both alone as well as in conjunction with the other scenario: $P(S_1)$ occupies the space in which S_1 occurs together with S_2, and $P(S_2)$ occupies the space in which S_2 occurs together with S_1. There is, however, only one space for $S_1 \,\&\, S_2$ as a combined scenario, and hence the deduction.

A joint occurrence of two (or more) events is not always factually possible. For example, a single toss of a coin can yield either heads or tails, but not both: that is, $P(S_1 \,\&\, S_2) = 0$. The coin's probability of landing on heads or, alternatively, on tails consequently equals 1 $(0.5 + 0.5 - 0)$. But in real-life situations, events often occur in conjunction with each other. For example, a medical patient's permanent disability may originate from his preexisting condition, from his doctor's malpractice, or from both. If so, then $P(S_1 \,\&\, S_2) > 0$.

A conjunctive occurrence of two events can also be perceived as a compound scenario in which one event (H) unfolds in the presence of another (E). The probability of any such scenario is called "conditional" because it does not attach unconditionally to a single event, H, but rather to event H given the presence, or occurrence, of E, which is denoted as $P(H|E)$.

This formulation allows me to present another core component of the mathematical probability system: Bayes' Theorem.[11] This theorem establishes that when I know the individual probabilities of E and H and the probability of E's occurrence in the presence of H, I can calculate the probability of H's occurrence in the presence of E. Application of the multiplication principle (the product rule) to the prospect of a joint occurrence of two events, E and H, yields $P(E \,\&\, H) = P(E) \times P(H|E)$. Under the same principle, the conjunctive probability of E and H, restated as $P(H\&E)$, also equals $P(H) \times P(E|H)$. This inversion sets up a probabilistically important equality: $P(E) \times P(H|E) = P(H) \times P(E|H)$.[12] Bayes' Theorem is derived from this equality: $P(H|E) = P(H) \times P(E|H) \div P(E)$.

My labeling of the two events as E and H is not accidental. Under the widely accepted terminology, H stands for a reasoner's *hypothesis*, while E stands for her *evidence*. Both E and H are events, but the reasoner is not considering those events individually. Rather, she is examining the extent to which evidence E confirms hypothesis H. A Bayesian formulation consequently separates between the probability of hypothesis H before the arrival of the evidence ($P(H)$); the general probability of the evidence's presence in the world ($P(E)$); and the probability of the evidence being present in cases in which hypothesis H materializes ($P(E|H)$). These three factors allow the reasoner to compute the posterior probability of her hypothesis: the probability of hypothesis H given evidence E. The reasoner must process every item of her evidence sequentially by applying this procedure. She must perform a Bayesian calculation time and time again until all of her evidence is taken into account. Each of those calculations will update the hypothesis's prior probability by transforming it into a new posterior probability. The posterior probability will become final after the reasoner had exhausted all of the available evidence.[13]

Notice the significance of the evidence-based multiplier, $P(E|H) \div P(E)$. This multiplier is called the "likelihood ratio" (Schum 1994, p. 218) or—as I prefer to call it—the "relevancy coefficient."[14] The relevancy coefficient measures the frequency with which E appears in cases featuring H, relative to the frequency of E's appearance in all possible cases. If $P(E|H) \div P(E) > 1$ (E's appearance in cases of H is more frequent than its general appearance), the probability of hypothesis H goes up. Formally: $P(H|E) > P(H)$, which means that evidence E confirms hypothesis H. On the other hand, when $P(E|H) \div P(E) < 1$ (E's appearance in cases of H is less frequent than its general appearance), the probability of hypothesis H goes down. Formally: $P(H|E) < P(H)$, which means that evidence E makes hypothesis H less probable (or disconfirms it). Finally, if $P(E|H) = P(E)$ (E's appearance in cases of H is as frequent as its general appearance), the presence of E does not influence the probability of H. This makes evidence E altogether irrelevant.[15]

To illustrate, consider a tax agency that uses internal fraud-risk criteria for auditing

---

[11] See Bayes (1763) for a modern statement of the theorem, see Cohen (1989, p. 68).

[12] Because of this inversion, some call Bayes' Theorem the "Inversion Theorem." See, e.g., Kneale (1949, p. 129).

[13] For a good explanation of this updating, see Schum (1994, pp. 215–22).

[14] Schum (1994, p. 219) associating the likelihood ratio with the "force of evidence".

[15] Cf. Lempert (1977, pp. 1025–27) offering similar formulation of relevancy coefficients.

firms.[16] By applying those criteria, the agency singles out for auditing one firm out of ten. This ratio is public knowledge. Firms do not know anything about the agency's criteria for auditing (nor does anyone else outside the agency). Under the information available to firms, their prior probability of being audited equals 0.1.

Now consider an individual firm whose reported expenses have doubled relative to past years. Does this evidence change the probability of being audited? The answer to this question depends on whether a steep increase in a firm's reported expenses appears more frequently in cases in which it was audited than in general. Assume that experienced accountants formed an opinion that increased expenses are three times more likely to appear in auditing situations than generally. This relevancy coefficient triples the prior probability of the firm's audit. The firm's posterior probability of being audited thus turns into 0.3.

But how do we know that these evidential effects are brought about by *causes* and are more than a mere correlation? We do not know it for sure, and I address this issue below in Section 2.2. My current goal, as I already mentioned, is quite narrow: in the present section, I only articulate the semantics and syntax of mathematical probability. Bayes' Theorem is part of those semantics and syntax: it tells us how to conceptualize our epistemic situations by using mathematical language. However, as I demonstrate below in Section 2.2, the theorem itself provides no instructions on how to understand causes and effects of the outside world and relate them to each other.

Mathematical language creates a uniform conceptual framework for all probability assessments that rely on instantial multiplicity or frequency of events. For those who base their estimates of probability on events' frequency, this language is indispensable.[17] This language is also necessary for formulating probability assessments on the basis of propensity—a disposition of a given factual setup to produce a particular outcome over a series of cases or experiments.[18] Finally, people basing their decisions upon intuited or "subjective" probabilities[19] might also benefit from using the mathematical language. This language introduces conceptual precision and coherence into a reasoner's conversion of her experience-based beliefs into numbers. Those numbers must more or less correspond to the reasoner's empirical situation. A mismatch between the numbers and empirical reality will produce a bad decision (Cohen 1989, p. 60).

Proper use of the mathematical language, however, does not guarantee that a person's probability assessments will be accurate. This language only helps a person conceptualize her raw information in numerical terms and communicate it to other people. Before using this language, a person must properly perceive and understand this information. This basic cognitional task is antecedent to a person's mathematical assessment of probability.

---

16 A good real-world example of this practice is the secret "Discriminant Index Function" (DIF), used by the IRS in selecting taxpayers for audits. See, e.g., *Gillin v. Internal Revenue Serv.* (980 F.2d 819, 822 1st Cir. [1992]) "The IRS closely guards information concerning its DIF scoring methodology because knowledge of the technique would enable an unscrupulous taxpayer to manipulate his return to obtain a lower DIF score and reduce the probability of an audit."; Lawsky (2009, pp. 1068–70) describing the DIF method used by the IRS.

17 See Cohen (1989, pp. 47–48) explaining frequency as a rate of relevant instances.

18 Cohen (1989, pp. 53–58) explaining propensity as a rate of relevant instances.

19 Cohen (1989, pp. 58–70) explaining subjective probability in terms of reasoners' betting odds.

Bayes' Theorem and other mathematical rules of probability do not tell people how to go about this task.

Proper use of mathematical probability therefore can only guarantee a gambling kind of accuracy: accuracy in ascribing probability estimates to perceived generalities, as opposed to individual events. If so, granted that a person properly conceptualizes her experiences in mathematical language, will her probability assessments be accurate if she commits no mathematical errors in making those assessments? This question is fundamental to the entire probability theory, and the answer to it depends on what "accurate" means. The mathematical system offers reasoners only one sort of guarantee. Absent information about relevant causes and effects, a reasoner will do well to follow that system, which would then enable her to achieve the maximal level of accuracy. Failure to follow that system will lead the reasoner astray.

This virtue of mathematical probability is best illustrated by a gambling scenario known as "Dutch Book." Consider a gambler who accepts two $100 bets at odds of 1 to 2 that a particular tennis player will win and, respectively, lose her upcoming match. This combination of bets is fundamentally irrational. Should the player win the match, the gambler would win $100 on the first bet, but would lose $200 on the second; and in the event the player loses the match, the gambler would lose $200 on the first bet and win only $100 on the second. Hence, the gambler is sure to lose $100.

This outcome has a simple explanation: the gambler ascribed an identical probability (0.667)[20] to factual propositions that negate one another, which was a bad idea. If the gambler's acceptance of the first bet were a good decision, the player's probability of winning the match would then be 0.667, as estimated by the gambler. Under that probability, however, the gambler could not rationally accept the second bet, which assumed that the player had a 0.667 probability of losing the match. Given that the gambler was right to accept the first bet, this probability could only be 0.333 (1 – 0.667). Any other probability assessment in placing bets would make the gambler lose his money.

Based on this insight, Frank Ramsey and Bruno de Finetti have demonstrated (independently of each other) that failure to follow the rules of mathematical probability engenders irrational decisions (Cohen 1989, pp. 60–61). This demonstration, however, holds true only in gambling situations in which decision-makers have no information about causes and effects that determine the course of specific events. Economists do not pay much attention to this limitation (Stein 2011, pp. 223–34), and I now turn to discuss it.

2.2 The Epistemics of Mathematical Probability

John Stuart Mill sharply criticized the use of instantial multiplicity as a basis for inference (Mill [1843] 1980, pp. 549–53). He described it as "the natural Induction of uninquiring minds, the induction of the ancients, which proceeds *per enumerationem simplicem*: 'This, that, and the other A are B, I cannot think of any A which is not B, therefore every A is B'" (Mill [1843] 1980, p. 549).

This sentence succinctly identifies the core problem of the mathematical probability

[20] This probability reflects the gambler's belief that the player wins two matches out of three under given conditions, which makes him accept the bet at the odds of 1 to 2.

system: this system, says Mill, is epistemologically fragile, if not empty. The system's mathematical rules only instruct the reasoner on how to convert her information into cardinal numbers. These rules have no epistemic ambition. They do not tell the reasoner what counts as information upon which she ought to rely. This task is undertaken by the system's rules of inference that are not as rigorous and intuitive as Boolean algebra. I examine those rules of inference in the paragraphs ahead.

One of those rules holds that any scenario not completely eliminated by existing evidence is a factual possibility that must occupy some of the probability space. The reasoner must consequently assign some probability to any such scenario, and this probability must be greater than zero. I call this rule "the uncertainty principle."

The second rule—"the principle of indifference"—is a direct consequence of the first. This rule determines the epistemic implications of the unavailable information for the reasoner's probability decision. The rule postulates that unavailable information is not slanted in any direction, meaning that the reasoner has no reasons for considering one unevidenced scenario as more probable than another unevidenced scenario (Cohen 1989, pp. 43–44). In other words, the reasoner ought to be epistemically indifferent between those scenarios, and this indifference makes the unevidenced scenarios equally probable.

The third rule logically derives from the second. It presumes that statistical distributions are extendible. To follow Mill's formulation, if 70% of events exhibiting feature *A* exhibit feature *B* as well, then presumptively any future occurrence of *A* has a 70% chance of occurring together with *B*. I call this rule "the extendibility presumption." This presumption is tentative and defeasible: new information showing, for example, that *B* might be brought about by *C*—a causal factor unassociated with *A*—would render it inapplicable. Absent such information, however, the extendibility presumption applies with full force. The presumption's mechanism relies on the indifference principle as well. This principle treats all indistinguishable occurrences of *A*, past and future, as equivalents. The same principle marks any missing information that could identify *B*'s causal origins as unslanted. The reasoner consequently must treat this unknown information as equally likely to both increase and decrease the rate of *B*'s appearance in cases of *A*. Every future occurrence of *A* thus becomes statistically identical to *A*'s past occurrences that exhibited *B* at a 70% rate.

The uncertainty principle seems epistemologically innocuous, but this appearance is misleading. Any factual scenario that existing evidence does not completely rule out must, indeed, be considered possible. This scenario therefore must have *some* probability on a 0–1 scale. All of this is undoubtedly correct. The uncertainty principle, however, also suggests that the reasoner *can* assign concrete probabilities to such unevidenced scenarios. This "can" is epistemologically unwarranted because the reasoner does not know those probabilities. Any of her probability estimates will be pure guesswork: a creation of knowledge from ignorance.

The principle of indifference is a pillar of the entire system of mathematical probability.[21] It stabilizes the reasoner's information in order to make it amenable to mathematical calcu-

[21] See Keynes (1921, pp. 41–42) describing the indifference principle as essential for establishing equally probable possibilities—a preliminary condition for all mathematical assessments of probability.

lus.[22] The principle's information-stabilizing method is best presented in Bayesian terms. Take a reasoner who considered all available information and determined the probability of the relevant scenario, $P(S)$. The reasoner knows that her information is incomplete and turns to estimating the implications of the unavailable information (U). The reasoner tries to figure out whether this unavailable information could change her initial probability estimate, $P(S)$. In formal terms, the reasoner needs to determine $P(S|U)$. Under Bayes' Theorem, this probability equals $P(S) \times [P(U|S) \div P(U)]$. With the prior probability, $P(S)$, already known, the reasoner needs to determine the relevancy coefficient, $P(U|S) \div P(U)$. To this end, she needs to obtain two probabilities: the probability of U's appearance in general and the probability of U's appearance in cases of S. Because the reasoner has no information upon which to make that determination, the indifference principle tells her to assume that U is equally likely to confirm and to disconfirm S: $P(U|S) = P(U)$. The relevancy coefficient consequently equals 1, and the reasoner's prior probability, $P(S)$, remains unchanged. The indifference principle essentially instructs the reasoner to deem missing information altogether irrelevant to her decision.

This instruction is epistemologically invalid. The reasoner can treat unavailable information as irrelevant to her decision only if she has no reason to believe that it might be relevant (Keynes 1921, pp. 55–56). Whether those reasons are present or absent depends on the reasoner's known information. When this information indicates that the unavailable information might be relevant, $P(U|S)$ and $P(U)$ can no longer be considered equal to each other. The indifference principle consequently becomes inapplicable. On the other hand, when the known information indicates that the unavailable information is irrelevant to the reasoner's decision, something else happens. The known information establishes that $P(U|S)$ actually equals $P(U)$. The proven, as opposed to postulated, equality between $P(U|S)$ and $P(U)$ makes the indifference principle redundant. From the epistemological point of view, therefore, there are no circumstances under which this principle can ever become applicable.[23]

The indifference principle thus does not merely purport to manage unavailable information. Instead, it forces itself on the available information by requiring the reasoner to interpret that information in a particular way. Effectively, the principle instructs the reasoner to proceed on the assumption that all the facts necessary for her probability assessment are specified in the available information. This artificially created informational closure sharply contrasts with the causative probabilistic reasoning that I discuss in Section 3.[24]

From an epistemological standpoint, the extendibility presumption is an equally problematic device. This presumption bypasses the question of causation, which makes

---

22 As Keynes explains, "In order that numerical measurement may be possible, we must be given a number of *equally* probable alternatives" (Keynes 1921, p. 41).

23 See Cohen (1989, pp. 45–46) showing that the indifference principle is either circular or redundant; Keynes (1921, pp. 45–47) demonstrating that the indifference principle is arbitrary and epistemologically unsustainable.

24 See Cohen (1979, p. 389), "Baconian [causative] probability-functions . . . grade probabilification . . . *by the extent to which* all relevant facts are specified in the evidence."

it epistemologically deficient.[25] As Mill's quote suggests, 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 (Cohen 1986, p. 177). This demonstration involves proof that *B* is or tends to be uniformly present in cases of *A* for reasons that remain the same in all cases (Cohen 1986, pp. 177–79). Those invariant reasons make the uniformity law-bound (Cohen 1986, p. 179). Their absence, in contrast, indicates that *B*'s presence in cases of *A* is possibly accidental. The observed uniformity consequently becomes non-extendible. Decision-makers who choose to rely on this uniformity will either systematically err or arrive at correct probability assessments by sheer accident. They will never base those assessments upon knowledge.[26]

To illustrate, consider again the basic factual setup of my tax-audit example: the tax agency audits one firm out of ten. Assuming that no other information is available, will it be plausible to estimate that each firm's probability of being audited equals 0.1? This estimate's plausibility depends on whether the "one-to-ten" distribution is extendible. This distribution could be extendible if the agency were to make its audit decisions by some randomized procedure, such as a draw. This randomization would then give every firm an equal chance of being audited by the agency. The agency, however, does not select audited firms by a draw. Instead, it applies its secret fraud-risk criteria. This fact makes the observed distribution of audits non-extendible. Consequently, the 0.1 estimate of a firm's probability of being audited is completely implausible. Relying on it would be a serious mistake.[27]

To rebut this critique, adherents of mathematical probability might invoke the long-run argument, mistakenly (but commonly) grounded upon Bernoulli's law of large numbers (Bernoulli [1713] 2006, pp. 315–40).[28] This argument concedes that the 0.1 estimate of a firm's probability of being audited is not a reliable predictor of any specific auditing event. The argument, however, holds that repeat-players—firms that file tax reports every year—should rely on this estimate because at some point it will transform into a real audit. With some firms, it will happen sooner than with others, but eventually the agency will audit every firm.

This argument recommends that every person perceive her epistemic state of uncertainty as an actual experience of a series of stochastic events that can take her life in any

---

[25] Another problem with extendibility is its dependence on a reference class—a statistical generalization that can be gerrymandered in numerous ways (Allen and Pardo 2007, pp. 111–14).

[26] For classic accounts of why accidentally true beliefs do not constitute knowledge, see Gettier (1963), which explains that accidentally acquired justification for a true belief is not knowledge; and Goldman (1967) attesting that a knower's true belief must be induced by the belief's truth. See also Nozick (1993, pp. 64–100) defining knowledge as a true belief supported by the knower's truth-tracking reasons.

[27] Taxpayers' responses to an increase in the general probability of audit are difficult to measure. For one such attempt, see Slemrod, Blumenthal and Christian (2001, p. 465), which finds that audit rates are positively correlated with reported income of low-income and middle-income taxpayers and are negatively correlated with reported income of high-income taxpayers.

[28] For a superb account of the law's intellectual history, see Hacking (1990, pp. 95–104).

direction. This recommendation fills every informational gap with God playing dice. However, neither God nor the tax agency will actually throw a die to identify firms that require an audit. Whether a particular firm will be audited will be determined by causal forces, namely, the tax officers who will apply the agency's fraud-risk criteria to what they know about each firm. Each firm therefore should rely on its best estimate of how those officers will evaluate its tax return. If, instead of relying on this estimate, a firm chooses to base its actions on the 10% chance of being audited, it will sooner or later find itself on the losing side.[29] This firm will either take wasteful precautions against liability for tax evasion or expose itself to that liability by acting recklessly.[30]

---

[29] This point was famously made by Samuelson (1963).

[30] To mitigate this problem, statisticians often use "confidence intervals." See, e.g., Wonnacott and Wonnacott (1990, pp. 253–86). A confidence interval is essentially a second-order probability: an estimate of the chances that the reasoner's event-related (first-order) probability is accurate. Conventionally, those chances must not go below 95%—a confidence level that promises that the reasoner's estimate of the event-related probability will be accurate in 95 cases out of 100 (Wonnacott and Wonnacott 1990, pp. 254–55). The reasoner must conceptualize her estimate of the event-related probability not as a fixed figure, but rather—more realistically—as an average probability deriving from a sample of probabilities attaching to factual setups similar to hers. The reasoner should expand her sample of setups by relying on her experience or by conducting a series of controlled observations. If she obtains a sufficiently large sample, the setups' probabilities will form a "normal" bell-shaped distribution curve. Subsequently, in order to obtain a 95% confidence level in her estimate of the probability, the reasoner must eliminate the curve's extremes and derive the estimate from the representative middle. Technically, she must shorten the distribution curve by trimming away 2.5% from each tail. This trimming will compress the reasoner's information and narrow the range of probabilities in her sample. The average probability calculated in this way will then have a high degree of accuracy. The chances that it will require revision in the future as a result of the arrival of new information are relatively low. This feature will make the probability estimate resilient or, as some call it, robust or invariant. See Logue (1995, pp. 78–95) associating strength of probability estimates with resiliency; Nozick (2001, pp. 17–19, 79–87) associating strength of probability estimates with their invariance across cases. The 95% confidence-interval requirement undeniably improves the quality of probabilistic assessments. The fact that those assessments stay invariant across many instances makes them dependable (Cohen 1989, p. 118). This improvement, however, does not resolve the deep epistemological problem identified in this section. Resilience of a probability estimate only indicates that the estimate is statistically stable. For example, a resilient probability of 0.7 can only identify the number of cases—70 out of 100—in which the underlying event will actually occur. This assurance, however, does not determine the applicability of the 0.7 probability to individual events. Whether this (or other) probability attaches to an individual event does not depend on the availability of this assurance; rather, it depends on the operation of the indifference principle and the extendibility presumption. These inferential rules apply to an individual event in the absence of information accounting for the difference between the cases in which the event occurs and the cases in which it does not occur. The reasoner will thus always make an epistemically unwarranted assumption that the unavailable information is not slanted in any direction. The mathematical system may try to adopt a more demanding informational criterion: one that differentiates between probability estimates on the basis of their epistemic weights (Keynes 1921, pp. 71–77). For contemporary analyses of Keynes's "weight" criterion, see Cohen (1989, pp. 102–109); Schum (1994, pp. 251–57); Stein (2005, pp. 80–91); Cohen (1985). Charles Peirce also endorsed this criterion when he observed that "to express the proper state of our belief, not *one* number but *two* are requisite, the first depending on the inferred probability, the second on the amount of knowledge on which that probability is based" (Peirce 1872–1878). Under this criterion, the weight of a probability estimate will be determined by the comprehensiveness of what the reasoner does and does not know about her case (Keynes 1921, pp. 71, 77). The decisional

3. CAUSAL PROBABILITY AND COMMON SENSE

Consider the following scenario:

> Peter undergoes a brain scan by MRI, and the scan is analyzed by a radiologist. The radiologist tells Peter that the lump that appears on the scan is benign to the best of her knowledge. She clarifies that she visually examined every part of Peter's brain and found no signs of malignancy. Peter asks the radiologist to translate the "best of her knowledge" into numbers, and the radiologist explains that 90% of the patients with similar-looking lumps have no cancer and that indications of malignancy are accidentally missed in 10% of the cases. The radiologist also tells Peter that only complicated brain surgery and a biopsy can determine with certainty whether he actually has cancer. According to the radiologist, this surgery involves a 15% risk of severe brain damage; in the remaining 85% of the cases, it successfully removes the lump and the patient recovers. Peter's primary care physician subsequently informs him that MRI machines have varying dependability. Specifically, he tells Peter that about 10% of those machines fail to reproduce images of small-size malignancies in the brain.

Under the mathematical system, Peter's probability of not having cancer equals 0.81. This number aggregates two probabilities of 0.9: the probability of correctness that attaches to the radiologist's diagnosis and the machine's probability of properly reproducing the image of Peter's brain. Peter's probability of having cancer consequently equals 0.19 (1 – 0.81).[31] This probability is greater than the 0.15 probability of sustaining severe brain damage from the surgery. Should Peter opt for the surgery?

Under the mathematical system, he should. The fatalities to which the two probabilities attach are roughly identical. If so, Peter should choose the course of action that reduces the fatality's probability. Under the mathematical system of probability, this choice will improve Peter's welfare (by 4% of the value of his undamaged brain).

Common sense, however, would advise Peter to rely on the causative probability instead. Specifically, it would tell Peter to rely on the radiologist's negative diagnosis and pay little or no attention to the background statistics. The radiologist's diagnosis is the only empirically-based causal account that concerns Peter's *individual* condition.[32] The radiologist informs Peter about what she saw and what did not see in his brain.[33] This diagnosis is the only information compatible with the causal nature of Peter's

synergy between probability and weight will create a serious problem of incommensurability. Consider a reasoner who faces a high but not weighty probability, on the one hand, and a weighty but low probability, on the other hand. Which of the two probabilities is more dependable than the other? This question does not have a readily available answer. There is simply no metric by which to compare the two sets of probabilities. This problem may not be insurmountable, but why tolerate it in the first place? Why try hard to undo the damage caused by the mathematical system's epistemological outlaws, instead of barring them? Sections 3 and 4 below respond to this question.

[31] This calculation applies the negation rule. The same probability can be calculated by aggregating Peter's 10% chance of having a small malignancy missed by the MRI machine with his 10% chance of being one of the radiologist's false negatives. Peter's probability of falling into *either* of these misfortunes equals (0.1 + 0.1) − (0.1 × 0.1) = 0.19. This calculation follows the disjunction rule.

[32] Cf. Cohen (1980) arguing that patient-specific diagnoses are superior to statistical ones.

[33] See, e.g., Mavroforakis et al. (2005) specifying malignancy and benignancy indicators that a radiologist should evaluate qualitatively in each patient and developing a quantitative tool to make those evaluations more robust.

physical environment. The general statistic extrapolated from the radiologist's and the MRI machine's history of errors is incompatible with this environment. This statistical information identifies no causal factors relevant to Peter's brain.

This common sense (that gets philosophical support from Francis Bacon [Bacon 1889] and John Stuart Mill [Mill [1843] 1980, pp. 549–53]) is impeccable. Peter should rely on the radiologist's diagnosis of his brain. Peter will make a serious and potentially fatal mistake if he chooses to undergo the brain surgery instead. Evidence that the radiologist erred in the past in ten diagnoses out of 100 reduces the general reliability of her diagnoses. This evidence, however, is causally irrelevant to the question of whether Peter has cancer. Whether Peter has cancer is a matter of empirical fact that the radiologist tried to ascertain. Her ascertainment of this fact relied on a series of patient-specific observations and medical science. While doing her job, the radiologist does not proceed stochastically by randomly distributing ten false-negative diagnoses across one hundred patients. Rather, she does her best for each and every patient, but, unfortunately, fails to identify cancer in 10 patients out of 100.

These errors had patient-specific or scan-specific causes: invisible malignancies, malfunctioning MRI machines, accidental oversights, and so forth. Those causes are unidentifiable, which means that Peter may still find himself among the afflicted patients. As an empirical matter, however, the unknown status of those causes does not equalize the chances of being misdiagnosed for each and every patient. Consequently, Peter has no empirical basis to discount the credibility of the radiologist's diagnosis of his brain by 10%.[34] This diagnosis is not completely certain, but it gives Peter qualitatively the best information that he can depend upon. This information is qualitatively the best because it is supported by an established causal theory: the radiologist's application of medical science to what she saw in Peter's brain. By contrast, no causal theory can ever support the view that the radiologist's patients are equally likely to be misdiagnosed as cancer-free.[35]

With this in mind, consider how ordinary people reason about their daily affairs. People are born into the world of causes and effects. Their daily affairs encompass events and phenomena that bring about other events and phenomena. As people accumulate their experiences and education, they internalize the idea of causation and the corresponding belief that things always happen for a reason and never without a reason.[36] Causal mechanisms underlying events and phenomena that people experience in their lives are not always known, but they are always present in the world. This causal understanding of the world drives most of ordinary people's decisions. These decisions therefore virtually always focus on some discrete, individual event and its underlying cause, as opposed to general distributions of similar-looking events. The same goes for generalizations that

[34] The same holds true for a possible malfunctioning of the MRI machine that scanned Peter's brain. There is no reason to believe that the risk of malfunction is distributed evenly across all machines and patients.

[35] Error statistics are not immaterial: if many (say, 30%) of the radiologist's diagnoses were false, Peter would have a good reason to doubt her credibility. This factor, however, would still be causatively irrelevant to whether he actually has cancer. Under these circumstances, Peter would have to find a credible specialist or endure the uncertainty. Cf. Thomson (1984, pp. 127–33) distinguishing between "external" evidence that derives from naked statistics and "internal" case-specific evidence that fits into a causal generalization.

[36] See generally Lagnado (2011).

ordinary people use in their decisions. These generalizations explain the world as governed by causal laws. They are akin to law-like generalities investigated by modern scientists.[37]

For that reason, when a person decides under conditions of uncertainty whether a certain event will (or did) occur, she articulates the available scenarios and selects the most plausible of those scenarios. More precisely, the person tries to figure out which of the available scenarios makes most sense in terms of coherence, consilience, causality, and evidential coverage (Allen and Stein 2013, pp. 567–71). This reasoning to the "best explanation" generally aligns with the common sense that people use in their daily affairs (Allen and Stein 2013, pp. 575–77).

This mode of reasoning rejects the indifference principle that animates mathematical probability. As I already explained, the indifference principle instructs reasoners to ignore the uncertainties in their evidence on the assumption that those uncertainties cancel each other out (Keynes 1921).[38] This assumption converts the reasoners' ignorance into the actual knowledge of probabilities, which it deems to be equal; and it has no epistemic warrant for that. Under the "best explanation" criterion, the decision-maker must select the best evidenced set of causes and effects, while rejecting all unevidenced hypotheses. She cannot assume that those hypotheses are equally probable—and thus cancel out—just because they are completely unevidenced. This epistemological injunction also does not allow the decision-maker to translate her reasons into mathematical fractions occupying a 0–1 scale. Because the decision-maker's information is incomplete, she has no epistemically justified reason to postulate that she knows the probabilities of all relevant scenarios. The decision-maker must consequently use words, rather than numbers, in evaluating the coherence, consilience, causal fit, and evidential coverage of competing scenarios.

To properly understand how ordinary people reason, one also needs to separate their "beliefs" from "acceptances," as recommended by philosophers of rationality.[39] Under this taxonomy, "acceptance" is a mentally active process that includes application of decisional rules to available information (Kahneman 2011, pp. 16–20). "Belief," by contrast, is a person's feeling, sensation, or hunch: an intellectually passive state of mind generated by unanalyzed experiences (Kahneman 2011, pp. 16–20).

Many of people's actions are driven by beliefs that people do not bother to reflect upon until it becomes necessary. For example, a person may form a belief that all medications sold by drugstores across the United States are as safe and as effective as advertised. Acting upon this unexamined belief is rational up to a point. For example, a person can rationally rely on this experience-based belief when she takes care of minor aches and discomforts. However, in serious health matters, a person will do well to discuss the

---

37 See generally Lange (1993) defining law-bound regularities as separate from accidental events.

38 See Keynes (1921, pp. 41–42) describing the indifference principle as essential for establishing equally probable possibilities—a preliminary condition for all mathematical assessments of probability; as Keynes explains, "In order that numerical measurement may be possible, we must be given a number of *equally* probable alternatives" Keynes (1921, p. 41, original emphasis); see Cohen (1989, 45–46) showing that the indifference principle is either circular or redundant; Keynes (1921, 45–47) demonstrating that the indifference principle is arbitrary and epistemologically unsustainable; see also Cohen (1979, p. 389) "Baconian [causative] probability-functions . . . grade probabilification . . . *by the extent to which* all relevant facts are specified in the evidence."

39 See Cohen (1992, pp. 1–27, 100–108) delineating the differences between "belief" and "acceptance".

pros and cons of every relevant medication with a qualified professional. Her decisions in such matters must rely upon rigorous and well-articulated criteria for assessing the medication's effects, as in my radiologist example. In other words, instead of simply relying on her beliefs, the person must form an "acceptance" based upon rules of reasoning.

Importantly, "belief" and "acceptance" are not analogs of what psychologists call "System 1" and "System 2" (Kahneman 2011, pp. 19–30). The "System 1/System 2" taxonomy only captures the intensity of a person's brainwork. To this end, it focuses on whether the person puts deliberative effort into her decisions (System 2) or decides quickly and unreflectively by using her intuition (System 1) (Kahneman 2011, pp. 1–30). By contrast, the belief–acceptance taxonomy captures the brainwork's normative content by separating the person's rule-free decisions (beliefs) from his rule-driven decisions (acceptances). System 1 and System 2 can, however, generate both beliefs and acceptances, depending on whether the person follows decisional rules—intuitively or reflectively. To be sure, a rule follower will use System 2 more often than System 1. Many people, however, also develop rule-driven instincts: drivers following the "two-second rule" to avoid colliding with a vehicle ahead of them are a good example of persons making rule-driven decisions that fall under System 1. On the other hand, some people may expend their deliberative efforts (System 2) on the formation of rule-free beliefs.

Behavioral economists systematically ignore these perfectly rational characteristics of ordinary people's reasoning. In Section 4 below, I evaluate the consequences of that omission. Before conducting that evaluation, I complete my discussion of causal probability by taking a closer look at its virtues. Specifically, I develop an analytical tool for separating cases that call for the application of mathematical probability from cases in which mathematical probability leads reasoners astray and where they will do well to use causal probability instead.

4. BOUNDED PROBABILISTIC RATIONALITY REVISITED[40]

In the following paragraphs I revisit the flagship experiments that helped behavioral economists to establish the bounded probabilistic rationality theory (BPR). My critique of BPR is twofold. First, I show that BPR and its supporting experiments suffer from insurmountable methodological problems. Subsequently, I demonstrate that BPR is flawed from the standpoint of conventional probability theory as well.

4.1 Belief vs. Acceptance

Behavioral experiments underlying the bounded rationality thesis uniformly miss the belief–acceptance distinction. People who participate in these experiments develop no rule-based acceptances, nor are they asked to form such acceptances by the experimenters. All they do is report their pre-analytical beliefs because this is what the experimenters ask them to do. People's rationality, however, can only be evaluated by reference to their

[40] This section draws on Stein (2013).

acceptances that apply rules of reasoning.[41] Identifying the criteria, or rules, that people apply in their evaluations of probability consequently becomes crucial.

Behavioral economists systematically fail to investigate people's acceptances, as distinguished from their beliefs. As I explain below, this omission undermines BPR. Failure to separate rule-driven acceptances from rule-free beliefs has also led behavioral economists to conflate people's cognitive performance with cognitive competence.[42] This conflation makes the resulting behavioral accounts deficient. The fact that a person systematically makes statistical errors in forming her beliefs does not establish that she would also commit those errors in forming her acceptances, in which case she would familiarize herself with and reflectively apply the requisite statistical rules. In fact, empirical studies of statistical education report considerable success of the various learning methods through which students acquire understandings of statistical inference.[43]

Behavioral economists' failure to separate beliefs from acceptances looms large in the "Linda Problem"—a celebrated experiment of Kahneman and Tversky (Kahneman 2011, pp. 156–58). Linda was described to participants as a 35-year-old woman, who was "single," "outspoken," "very bright," and deeply concerned with "issues of discrimination and social justice." Linda's college life included majoring in philosophy and participating in anti-nuclear demonstrations. Participants were asked to select Linda's occupation and social identity from the list provided by Kahneman and Tversky. "Bank teller" and "feminist bank teller" were among the options on that list. Most participants ranked Linda's being a "feminist bank teller" as more probable than Linda's simply being a "bank teller."

This assessment of probability defies mathematical logic. Linda's feminism was a probable, but still uncertain, fact. Her occupation as a bank teller was a merely probable, rather than certain, fact as well. The probability of each of those characteristics was somewhere between 0 and 1. Hence, the probability that these two characteristics would be present simultaneously must be lower than the probability that attached to each individual characteristic. Linda was more likely to have only the "bank teller," or only the "feminist," characteristic than to possess both characteristics at once. Assuming that the characteristics are mutually independent and that the probability of each characteristic is, say, 0.6, Linda's probability of being a feminist bank teller would equal 0.36. Remarkably, the Linda results were replicated with doctorate students at Stanford Business School.

To verify this important finding, Kahneman and Tversky conducted another experiment that featured a simple question: "Which alternative is more probable? Linda is a bank teller. Linda is a bank teller and is active in the feminist movement." Once again, the participants ranked the second joint-event scenario as more probable than the first single-event scenario.

Kahneman reports that after completing one such experiment, he asked the participants, "Do you realize that you have violated an elementary logical rule?" In response,

---

41 Stein (2013, p. 88) "Nor would [a belief] deserve praise or blame in the way that a responsible act of acceptance deserves it."

42 This conflation was first spotted by Cohen (1981, pp. 328–29).

43 See, e.g., Arnold et al. (2011) reporting success with teaching statistical inference to 14-year-old students with the help of hands-on physical simulations; Hall and Vance (2010), reporting success in teaching introductory statistics with the help of students' self-explanation and peer feedback.

a graduate student said "I thought you just asked for my opinion." Kahneman cites this response to illustrate the stickiness of people's probabilistic irrationality: the student who gave this response believed that her *opinion* on factual matters could defy mathematical logic.

The student's response, however, ought to have moved Kahneman in a different direction. What the student was actually saying was "Had I known that you were expecting me to give you not just my best hunch about Linda's job and social identity, but rather a rule-based evaluation of the relevant probabilities, my answer might have been different." The student, in other words, understood the experiment as asking her to express her belief, rather than articulate and apply her criteria for acceptances. In forming this belief, she felt free to rely on her common sense and experience rather than on statistical rules. Her reasoning aligned with that of scientists who begin their inquiries with intuitive beliefs that they subsequently accept or reject (Cohen 1989, pp. 89–90).

Similar to many other experiments carried out by behavioral economists, the Linda Problem could only elicit the beliefs that participants intuitively formed. Those beliefs do not reveal much about the participants' probabilistic rationality. Forming a rule-free belief, as opposed to a rule-driven acceptance, does not commit the believer to any specific reason, or rule, that she will follow in her other decisions.[44] Acceptances driven by rules of reasoning are different. Most medical patients, for example, would attest that having spinal surgery followed by a coronary bypass operation is riskier than undergoing spinal surgery alone. This attestation correctly applies the product rule for conjunctive probabilities to facts that the patient deeply cares about. Unsurprisingly, it expresses the patient's acceptance rather than belief.

As far as beliefs are concerned, the participants' prevalent reaction to Linda was far from irrational. Formation of a person's belief always calls in the experience that a person has accumulated throughout her life.[45] This experience cannot be artificially blocked by statistical rules, unless the person is expressly told to suppress all of her beliefs that do not conform to those rules and to base her decision on acceptance.[46] From the standpoint of an ordinary person's belief, the single-event scenario "Linda is a bank teller" was incomplete because bank tellers' work does not normally occupy their entire lives. The absence of information about Linda's social identity and afterwork engagements thus created a gap fillable by experience. Hence, it was entirely rational for participants to make an experience-based assumption that Linda must have *some* social identity or afterwork engagement. This assumption made the participants focus on the following question: is it more probable that, "Linda is a feminist bank teller" or that "Linda is a bank teller whose social identity and afterwork engagements are feminism free"?[47]

Based on Linda's background information, the participants were absolutely (and

---

[44] Cf. Schauer (1995) associating official reasons with commitment to apply similar reasons in future cases.

[45] See, e.g., Hume (1739) famously explaining "belief" as a consequence of the believer's "number of past impressions and conjunctions."

[46] Cf. Cohen (1991) arguing that jurors should suppress their beliefs and determine facts through "acceptance."

[47] Cf. Gigerenzer (2005, pp. 8–9) criticizing Linda and similar experiments for their reliance on a "content-blind" norm for rationality.

unsurprisingly) correct in forming a belief that ranked Linda's feminism above other afterwork engagements. In technical terms, Linda's probability of being a bank teller and a feminist, $P(T\&F)$, equaled $P(T) \times P(F)$. Correspondingly, Linda's probability of being a bank teller while having a non-feminist afterwork engagement, $P(T\&NF)$, equaled $P(T) \times P(NF)$. Under the factual setup that the participants were asked to consider, Linda was more likely to be a feminist than a non-feminist: $P(F) > P(NF)$. Hence, $P(T\&F) > P(T\&NF)$.

To preclude the formation of this rational belief, Kahneman and his associates ought to have asked the participants a simple question, suggested by Gerd Gigerenzer: "There are 100 persons who fit the description above (that is, Linda's). How many of them are: Bank tellers? Bank tellers and active in the feminist movement?" (Gigerenzer 2005, p. 10). This question would have elicited predominantly the statistically correct response (Gigerenzer 2005).

Kahneman's anticipated reply to this critique might fall along the following lines. The participants' real task was to cut through the "noise" (the statistically meaningless information) and see what the experimenters asked them to do. The participants, so goes the argument, ought to have noticed that their task was to compare the probabilities of a single and a compound, or conjunctive, event. Had the participants noticed that, they also would have noticed that Linda's probability of being a feminist bank teller was no different from the proverbial coin's probability of revealing *heads* in two successive throws. On a 0 to 1 scale, this probability equals $0.5 \times 0.5 = 0.25$.

The coin analogy, however, is untidy because Linda's social identity and afterwork engagement were not an unrigged coin. Linda's background information made her engagement in feminist causes the most probable afterwork scenario. Arguably, this scenario was more probable than the case in which Linda's work as a bank teller—surprisingly fulfilling or unduly exhaustive—represented everything she did in her life.

The upshot of my preceding discussion is straightforward. Studies of people's probabilistic decisions are not very fruitful when they focus on intuitive beliefs. Focusing on people's rule-driven acceptances in settings that call for statistical reasoning—as in my double-surgery example—could give Kahneman and other behavioral economists a much better sense of people's probabilistic rationality.

Behavioral economists, however, have chosen not to go along this route. Instead of adopting a simple all-statistics setup for their experiments, they mix statistical data with case-specific information. This informational mix can be found not only in the Linda Problem. Almost every experiment associated with the Kahneman and Tversky school of thought uses this mix, and there is a reason for that as well. Kahneman explains that causal associations corrupt people's decisions: people try to find causal connections where none exists, while irrationally discounting important statistical information (Kahneman 2011, pp. 74–78). This cognitive malfunction has shaped Kahneman's and his associates' experimental agenda. Kahneman and his associates seek to uncover how people's "causation illusion" drives them to ignore statistical data and depart from statistical reasoning. As I demonstrated in Section 3, however, there is nothing wrong in people's attempt to understand the outside world as a series of causes and effects. In fact, people will do well to rely on that understanding in most decisions they make during their lifetime.

4.2 BPR vs. Probability Theory

BPR also encounters difficulties in the realm of probability theory. Responsible for those difficulties is the statistical–causative mix of information on which behavioral economists often base their experiments. Consider one of Kahneman and Tversky's most famous experiments: the "Blue Cab Problem." Kahneman, Tversky, and their collaborators told their participants about a hit-and-run accident that occurred at night in a city in which 85% of cabs were blue and 15% were green (Kahneman 2011, pp. 166–70; Bar-Hillel 1980, pp. 211–12). They also told the participants that the hit-and-run victim filed a lawsuit against the companies operating those cabs—identified respectively as "Blue Cab" and "Green Cab"—and that an eyewitness testified in the ensuing trial that the cab that hit the victim was green. Another piece of information that the participants received concerned a rather unusual procedure that took place at this trial. The experimenters told the participants that "[t]he court tested the witness' ability to distinguish between Blue and Green cabs under nighttime visibility conditions [and] found that the witness was able to identify each color correctly about 80% of the time, but confused it with the other color about 20% of the time" (Bar-Hillel 1980, pp. 211–12). Based on this information, most participants in the experiments assessed the probability that a green cab hit the victim at 0.8, presumably because they believed this was the probability that the eyewitness's testimony was correct (Kahneman 2011, p. 167).

This assessment of probability aligned with the given credibility of the witness, but not with Bayes' Theorem.[48] The prior odds that the responsible cab was green as opposed to blue, $P(G)/P(B)$, equaled 0.15/0.85. To calculate the posterior odds, $P(G|W)/P(B|W)$, with W denoting the credibility of the witness, these odds had to be multiplied by the likelihood ratio. This ratio is equal to the odds attaching to the scenario in which the witness identified the cab's color correctly, rather than incorrectly: $P(W|G)/P(W|B)$. The posterior odds consequently equaled $(0.15 \times 0.8)/(0.85 \times 0.2)$—that is, 12/17. The probability that the victim's allegation against the Green Cab is true thus amounted to $12/(17 + 12)$ or 0.41—far below the "preponderance of the evidence standard" (> 0.5) that applies in civil litigation. The experiment thus seems to provide an elegant and robust demonstration of individuals' total neglect of base rates.

This and similar experimental vignettes have a serious flaw that I call unspecified causality. The experimenters did not tell the participants that the relative frequency of blue and green cabs' appearances on the streets of the city could somehow affect the witness's capacity to tell blue from green. This causal effect is quite unusual: an ordinary person can tell blue from green even when they see one green cab and many blue cabs.[49] The experimenters therefore ought to have told the participants that the witness's ability to distinguish between blue and green cabs might have been affected by the frequency with which those cabs appeared on the streets. Alternatively, the experimenters ought to have told the participants that in cases in which the witness failed to give the correct

48 For exposition and proof of Bayes' Theorem, see Stein (2011, pp. 211–13).

49 See Cohen (1986, p. 329) "[I]f the green cab company suddenly increased the size of its fleet relative to that of the blue company, the accuracy of the witness's vision would not be affected, and the credibility of his testimony would therefore remain precisely the same in any particular case of the relevant kind."

identification of the cab's color, he might have made this mistake randomly rather than for some specific reason (Stein 2011, pp. 253–55).

The experimenters, in other words, ought to have ruled the causality factor in or out. Instead, they allowed the participants to deal with the unspecified causality as they deemed fit, and the participants rendered an unsurprising—albeit not watertight—verdict that the distribution of cabs' colors in the city did not affect the witness's ability to tell blue from green. Absent a causal connection between these two factors, the errant cab's probability of being green as opposed to blue was indeed 0.8.

Unspecified causality is a serious flaw also because it makes the relevant reference class malleable.[50] To see how this malleability affected the Blue Cab Problem, factor in the preponderance requirement that a plaintiff in a civil suit needs to satisfy in order to win the case.[51] Under this requirement, the victim was certainly entitled to win her suit against Green Cab when the errant cab's probability of being green, given the testimony of the witness—$P(G|W)$—was greater than 0.5. The victim, however, was equally entitled to win the suit when the probability of the scenario in which the witness correctly identifies a green cab—$P(W|G)$—was greater than 0.5. The relevant reference class, in other words, could have been either the cab's color or the witness's accuracy.[52] The participants therefore could not be wrong in selecting the witness's accuracy as the relevant reference class. This perfectly rational choice allowed the participants to treat the probability of the witness's accuracy (0.8) as a dominant factor in their decision.

More fundamentally, the mix of statistical and causative information brings into consideration the normative openness of the "probability" concept (Stein 2011, pp. 200–204). As a normative matter, the Blue Cab Problem can be analyzed under two distinct analytical frameworks: mathematical (Pascalian) and causative (Baconian) (Stein 2011, pp. 253–56). The mathematical framework uses Bayes' Theorem, whose application gives the victim's case a 0.41 probability (if we ignore the unspecified causality and the reference-class problem). This probability represents the errant cab's chances of being green rather than blue, with a cab-identifying witness scoring 80 out of 100 on similar identifications in a city in which 85% of the cabs are blue and 15% are green.

The causative framework, on the other hand, yields an altogether different result, close to the mathematical probability of the witness's accuracy (0.8). Under this framework, which I explained in more detail in Section 3, an event's probability corresponds to the quantum and variety of the evidence that confirms the event's occurrence while eliminating rival scenarios (Stein 2011, pp. 243–46). This qualitative evidential criterion separates causative probability from the mathematical calculus of chances (Stein 2011, pp. 235–46). Under this criterion, the eyewitness's testimony that the errant cab was green was credible enough to rule out the "errant blue cab" scenario as causatively implausible. On the other hand, the distribution of blue and green cabs in the city had no proven effect on the eyewitness's capacity to tell blue from green. The eyewitness's testimony consequently overrode the cabs' distribution evidence and removed it from the fact-finding process.

---

50 For an outstanding analysis of reference-class malleability, see Allen and Pardo (2007, pp. 111–14).

51 See Stein (2005, pp. 143–48, 219–25) explaining the preponderance requirement and its underlying justifications.

52 This insight belongs to Owen (1987, p. 199).

This eliminative method (favored by Francis Bacon and John Stuart Mill [Stein 2011, pp. 204–206, 236–40]) allowed the participants to evaluate the probability of the victim's case at 0.8. This fact-finding method is not devoid of difficulties, but it is also far from being irrational (Stein 2011, pp. 236–40).

Contrary to Kahneman's view, the Blue Cab Problem and similar experiments do not establish that people's probability judgments are irrational.[53] These judgments *are* predominantly rational. The legal system need not do more than remedy people's informational shortfalls—not cognitive incapacities—by applying the conventional doctrines of foreseeability, disclosure, informed consent, unconscionability, and consumer protection.

4.3 Causation vs. Chance

Behavioral economists often criticize people for putting too much faith in causation (Kahneman 2011, pp. 74–78, 114–18). This criticism presupposes that incomplete causal indicators can only create an associative illusion of causation (Kahneman 2011). At the same time, behavioral economists believe that incomplete statistical indicators—the chances that surround us—are real and hence dependable (Kahneman 2011, pp. 71–78, 166–74).

This unexplained normative asymmetry is best illustrated by another milestone experiment of Kahneman and Tversky. Aimed at identifying the "representativeness" bias, the "Steve Problem" featured Steve, described to participants as "very shy and withdrawn, invariably helpful, but with little interest in people, or in the world of reality. A meek and tidy soul, he has a need for order and structure, and a passion for detail" (Kahneman 2011, p. 420). The experimenters asked the participants to choose Steve's most probable occupation from a list that included "farmer, salesman, airline pilot, librarian, [and] physician" (Kahneman 2011). According to Kahneman and Tversky, the participants used familiar (i.e., "representative") stereotypes to identify Steve as a likely librarian, while ignoring the fact that librarians are vastly outnumbered by farmers (Kahneman 2011).

Kahneman and Tversky assume that there was only one correct way to answer the question about Steve's job (Kahneman 2011, pp. 420–21). According to them, the participants had to find out the percentage of farmers, salesmen, airline pilots, librarians, and physicians in the general pool of working males. This percentage determined Steve's probability of being a farmer, a salesman, an airline pilot, a librarian, or a physician. Kahneman and Tversky believe that trying to identify Steve's profession through his personality traits is doomed to fail, as these traits are rather weak causal indicators of a person's professional identity. The general statistic representing an average working male's chances of having one of the above-mentioned professions was a far more dependable indicator. This indicator therefore ought to have trumped the uninformative individual traits. The participants' failure to notice this statistical indicator, and their

---

53 Cf. Gigerenzer (1996, p. 593) criticizing Kahneman and Tversky for testing people's ascriptions of probabilities to single events not amenable to such assessments.

consequent reliance on Steve's individual traits was a cognitive error (Kahneman 2011, pp. 420–21).

I posit that this experiment was poorly designed. Steve's personality traits did not make him a librarian, but they were certainly relevant to his choice of profession. If so, the participants should have been looking for a different, and more refined, statistic. Specifically, they should have been looking for the percentage of farmers, salesmen, airline pilots, librarians, and physicians in the general pool of working males who are shy, withdrawn and helpful, have meek and tidy souls and a passion for detail, and also need order and structure, while exhibiting little interest in people and the world of reality. Of course, this investigation would have been futile because general employment statistics do not single out the subcategory of working males formulated by Kahneman and Tversky. However, the fact that this investigation would have been futile does not make it inconsequential. Information revealing Steve's job preferences was material. Distribution of professions across working males generally was a rough and potentially misleading substitute for that information. This distribution was informative, but its evidential value did not outweigh the evidential value of Steve's personality traits. Kahneman and Tversky evidently think that it did, but this is just an opinion rather than empirical fact. People participating in the Steve Problem were entitled to use their opinions instead.

The Steve Problem's design incorporates unspecified causality. This feature opened two decisional routes for the experiment's participants. One could rationally estimate Steve's probability of being a librarian as a matter of chance. Alternatively, one could estimate this probability as a question of Steve's choice. Under the framework of chance, decision-makers would rely on the distribution of relevant professions across working males in general. Under the framework of choice, they would consider a probable bargaining equilibrium between Steve and prospective employers. This equilibrium solution would practically remove from the list the physician, the pilot, and the salesman. Arguably, as between being a farmer and being a librarian, Steve would choose to be a librarian. Finding a librarian position might be difficult—given the scarcity of such positions, relative to the many jobs available on a farm—but Steve could succeed in getting it.

Kahneman and Tversky disapprove of the participants' preference for the choice framework. Notwithstanding their disapproval, this preference is perfectly rational. The choice framework is not problem free, given the scarcity of case-specific information about Steve, but extrapolating Steve's probable occupation from the general pool of working males is equally problematic. Both modes of reasoning rely heavily on speculation, and there is consequently no way to tell which of them is epistemically preferable. Calling one of these modes of reasoning "rational" and another "irrational" is simply wrong.

Unspecified causality in an experiment's design *always* makes the relevant reference class malleable. Consider again Steve's case. Individuals participating in this experiment could perceive their task in two completely different ways. They could ask themselves whether Steve's personality traits separate him from the average working male. According to Kahneman and Tversky, this was the right question to ask. However, an alternative—and equally rational—way to define the reference class was to focus on a narrower category of working males who have Steve's characteristics. The relevant reference class, in other words, could be either of the following: (1) males, as distributed across different professions; or (2) professions, as distributed across different males. The first of these categories emphasizes chance, while the second centers on choice. There is no way to determine which of those categories is more dependable than the other as a basis for statistical inference. Kahneman and Tversky evidently prefer chance over choice. The participants in their experiments did the opposite. As for myself, I remain undecided.

Behavioral economists criticize people's reliance on case-specific knowledge as a "law of small numbers" (Kahneman 2011, pp. 109–18). This criticism is far removed from how most people—including judges and juries[54]—ascertain facts in their day-to-day lives. Behavioral economists' skepticism about case-specific knowledge also cannot be justified as a wholesale proposition, for it brushes aside a distinct mode of probabilistic reasoning, known as causative or Baconian reasoning (Stein 2011, pp. 204–206, 235–46). Behavioral economists' disregard of Baconian probability is perplexing. This mode of probabilistic reasoning is perfectly rational (Stein 2011), and it also could explain—and, indeed, justify—people's decisions that behavioral economists describe as erroneous.[55]

Under the causative system outlined in Section 3, a combination of credible case-specific evidence and experience can develop a single causal explanation for the relevant event that will override the competing statistical explanations (Stein 2011, pp. 235–46). This override is the essence of the Baconian elimination method (Stein 2011, pp. 204–206). For example, in the Blue Cab Problem, participants were entitled to assign overriding probative value to the witness's testimony that the errant cab was green. This testimony was not watertight, but it was credible and event specific. The event's causal impact on the witness's perceptive apparatus qualitatively differed from the city's cab-color statistics. This impact might have been epistemically superior to those statistics and therefore properly overrode them in the participants' minds.

This override was likely at work in the Steve Problem as well. There, participants used Steve's personality traits to eliminate from their list every profession that did not fit the stereotype associated with these traits. "Librarian" was the only item that survived this elimination procedure, which led the participants to estimate that Steve must be a librarian. Kahneman and Tversky correctly observe that this estimate was unfounded.[56] They are, however, too quick to denounce the participants' reasoning for failing to account for Steve's prior probability of being a farmer, as opposed to a librarian. Under the causative system of probability, the elimination method that the participants chose to use was valid. The participants simply did not have enough evidence for choosing the "librarian" over the "farmer." They would have had enough evidence for this assessment had they been informed, for example, that Steve is a connoisseur of literature. The addition of this information would have allowed the participants not to factor the statistical prevalence of farmers into their assessment.

---

[54] See Stein (2005, pp. 80–106) explaining case-specificity requirements in the law of evidence.

[55] See Cohen (1986, pp. 165–68) explaining why it is rational for people to rely on causative probabilities instead of naked statistics.

[56] They explained that "In the case of Steve . . . the fact that there are many more farmers than librarians in the population should enter into any reasonable estimate of the probability that Steve is a librarian rather than a farmer" (Kahneman 2011, p. 420).

5. CONCLUSION

Throughout their long history, humans have worked hard to tame chance.[57] They adapted to their uncertain physical and social environments by using the method of trial and error. This evolutionary process made humans reason about uncertain facts the way they do. Behavioral economists argue that humans' natural selection of their prevalent mode of reasoning wasn't wise. They censure this mode of reasoning for violating the canons of mathematical probability that a rational person must obey.

In this chapter, I have challenged both parts of this claim. Based on the insights from probability theory and the philosophy of induction, I have argued that a rational person need not apply mathematical probability in making decisions about individual causes and effects. Instead, she should be free to use common-sense reasoning that generally aligns with causative (Baconian) probability. I also have shown that behavioral experiments uniformly miss their target when they ask reasoners to extract probability from information that combines causal evidence with statistical data. Because it is perfectly rational for a person focusing on a specific event to prefer causal evidence to general statistics, those experiments establish no deviations from rational reasoning. Those experiments are also flawed in that they do not separate the reasoners' unreflective beliefs from rule-driven acceptances. The behavioral economists' claim that people are probabilistically challenged consequently remains unproven.

REFERENCES

Allen, Ronald J., and Michael S. Pardo. 2007. The Problematic Value of Mathematical Models of Evidence. *Journal of Legal Studies* 36: 10–40.

Allen, Ronald J., and Alex Stein. 2013. Evidence, Probability, and the Burden of Proof. *Arizona Law Review* 55: 557–602.

Arnold, Pip, Maxine Pfannkuch, Chris J. Wild, Matt Regan, and Stephanie Budgett. 2011. Enhancing Students' Inferential Reasoning: From Hands-On to "Movies". *Journal of Statistics Education* 19. http://www.amstat.org/publications/jse/v19n2/pfannkuch.pdf (last visited May 13, 2015).

Bacon, Francis. 1889. Novum Organum. P. 191 in *Bacon's Novum Organum*, 2nd edn, edited by Thomas Fowler. Oxford: Clarendon Press.

Bar-Grill, Oren, and Rebecca Stone. 2009. Mobile Misperceptions. *Harvard Journal of Law & Technology* 23: 49–118.

Bar-Grill, Oren, and Elizabeth Warren. 2008. Making Credit Safer. *University of Pennsylvania Law Review* 157: 1–101.

Bar-Hillel, Maya. 1980. The Base-Rate Fallacy in Probability Judgments. *Acta Psychologica* 44: 211–33.

Bayes, Thomas. 1763. An Essay Towards Solving a Problem in the Doctrine of Chances. http://www.stat.ucla.edu/history/essay.pdf (last visited Jan. 25, 2014).

Ben-Shahar, Omri, and Carl E. Schneider. 2011. The Failure of Mandated Disclosure. *University of Pennsylvania Law Review* 159: 647–749.

Bernoulli, Jacob. [1713] 2006. *The Art of Conjecturing*. Trans. Edith Dudley Sylla. Baltimore, MD: Johns Hopkins University Press.

Cohen, L. Jonathan. 1979. On the Psychology of Prediction: Whose Is the Fallacy? *Cognition* 7: 385–407.

Cohen, L. Jonathan. 1980. Bayesianism Versus Baconianism in the Evaluation of Medical Diagnoses. *British Journal for the Philosophy of Science* 31: 45–62.

Cohen, L. Jonathan. 1981. Can Human Irrationality Be Experimentally Demonstrated? *Behavioral and Brain Sciences* 4: 317–70.

---

[57] See generally Hacking (1990).

Cohen, L. Jonathan. 1985. Twelve Questions About Keynes's Concept of Weight. *British Journal for the Philosophy of Science* 37: 263–78.
Cohen, L. Jonathan. 1986. *The Dialogue of Reason: An Analysis of Analytical Philosophy*. New York, NY: Oxford University Press.
Cohen, L. Jonathan. 1989. *An Introduction to the Philosophy of Induction and Probability*. New York, NY: Oxford University Press.
Cohen, L. Jonathan. 1991. Should a Jury Say What It Believes or What It Accepts? *Cardozo Law Review* 13: 465–83.
Cohen, L. Jonathan. 1992. *An Essay on Belief and Acceptance*. Oxford: Clarendon Press.
Eisenberg, Melvin Aron. 1995. The Limits of Cognition and the Limits of Contract. *Stanford Law Review* 47: 211–59.
Gettier, Edmund L. 1963. Is Justified True Belief Knowledge? *Analysis* 23: 121–23.
Gigerenzer, Gerd. 1996. On Narrow Norms and Vague Heuristics: A Reply to Kahneman and Tversky (1996). *Psychological Review* 103: 592–98.
Gigerenzer, Gerd. 2005. I Think, Therefore I Err. *Social Research* 72: 195–218.
Goldman, Alvin I. 1967. A Causal Theory of Knowing. *Journal of Philosophy* 64: 357–72.
Hacking, Ian. 1990. *The Taming of Chance*. Cambridge: Cambridge University Press.
Hall, Simin, and Eric A. Vance. 2010. Improving Self-Efficacy in Statistics: Role of Self-Explanation & Feedback. *Journal of Statistics Education* 18. https://ww2.amstat.org/publications/jse/v18n3/hall.pdf (last visited May 13, 2015).
Hume, David. 1739. A Treatise of Human Nature. Pp. 1–102 in Vol. 1 of *A Treatise on Human Nature*, edited by David Fate Norton and Mary J. Norton. Oxford: Oxford University Press.
Jolls, Christine, Cass R. Sunstein, and Richard Thaler. 1998. A Behavioral Approach to Law and Economics. *Stanford Law Review* 50: 1471–550.
Kahan, Dan M. 2010. The Economics—Conventional, Behavioral, and Political—of "Subsequent Remedial Measures" Evidence. *Columbia Law Review* 110: 1616–53.
Kahneman, Daniel. 2011. *Thinking, Fast and Slow*. New York, NY: Farrar, Straus and Giroux.
Keynes, John Maynard.1921. *A Treatise on Probability*. London: Macmillan and Co.
Kneale, William. 1949. *Probability and Induction*. Oxford: Clarendon Press.
Kolmogorov, A. N. 1956. *Foundations of the Theory of Probability*. 2nd English edn. Translation edited by Nathan Morrison. New York, NY: Chelsea Publishing Company.
Lagnado, David A. 2011. Causal Thinking. Pp. 129–46 in *Causality in the Sciences* edited by Phyllis McKay-Illari, Federica Russo, and Jon Williamson. New York, NY: Oxford University Press.
Lange, Marc. 1993. Lawlikeness. *Noûs* 27: 1–27.
Lawsky, Sarah B. 2009. Probably? Understanding Tax Law's Uncertainty. *University of Pennsylvania Law Review* 157: 1017–74.
Lempert, Richard O. 1977. Modeling Relevance. *Michigan Law Review* 75: 1021–57.
Logue, James. 1995. *Projective Probability*. Oxford: Clarendon Press.
Mavroforakis, Michael, Harris Georgiou, Nikos Dimitropoulos, Dionisis Cavuras, and S. Theodoridis. 2005. Significance Analysis of Qualitative Mammographic Features, Using Linear Classifiers, Neural Networks and Support Vector Machines. *European Journal of Radiology* 54: 80–89.
Mill, John Stuart. [1843] 1980. *A System of Logic, Ratiocinative and Inductive: Being a Connected View of the Principles of Evidence and the Methods of Scientific Investigation*. 8th edn. New York: Harper.
Nozick, Robert. 1993. *The Nature of Rationality*. Princeton, NJ: Princeton University Press.
Nozick, Robert. 2001. *Invariances: The Structure of the Objective World*. Cambridge, MA: Harvard University Press.
Owen, David. 1987. Hume Versus Price on Miracles and Prior Probabilities: Testimony and Bayesian Calculation. *Philosophical Quarterly* 37: 187–202.
Oxford English Dictionary, The. 1989. 2nd edn. New York, NY: Oxford University Press.
Peirce, Charles Sanders. 1872–1878. The Probability of Induction. P. 295 in vol. 3 of *Writings of Charles S. Peirce*, *A Chronological Edition*, edited by Christian J. W. Kloesel. Bloomington, IN: Indiana University Press.
Rachlinski, Jeffrey J. 1998. A Positive Psychological Theory of Judging in Hindsight. *University of Chicago Law Review* 65: 571–625.
Samuelson, P. A. 1963. Risk and Uncertainty: A Fallacy of Large Numbers. *Scientia* 98: 108–13.
Schauer, Frederick. 1995. Giving Reasons. *Stanford Law Review* 47: 633–59.
Schum, David A. 1994. *The Evidential Foundations of Probabilistic Reasoning*. New York, NY: John Wiley & Sons, Inc.
Slemrod, Joel, Marsha Blumenthal, and Charles Christian. 2001.Taxpayer Response to an Increased Probability of Audit: Evidence from a Controlled Experiment in Minnesota. *Journal of Public Economics* 79: 455–83.
Stein, Alex. 2005. *Foundations of Evidence Law*. New York, NY: Oxford University Press.

Stein, Alex. 2011. The Flawed Probabilistic Foundation of Law & Economics. *Northwestern University Law Review* 105: 199–260.
Stein, Alex. 2013. Book Review: Are People Probabilistically Challenged? *Michigan Law Review* 111: 855–75.
Sunstein, Cass R. 1986. Legal Interference with Private Preferences. *University of Chicago Law Review* 53: 1129–74.
Thaler, Richard H., and Cass R. Sunstein. 2008. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. New York, NY: Penguin Books.
Thomson, Judith Jarvis. 1984. Remarks on Causation and Liability. *Philosophy and Public Affairs* 13: 101–33.
Vandenbergh, Michael, Amanda R. Carrico, and Lisa Schultz Bressman. 2011. Regulation in the Behavioral Era. *Minnesota Law Review* 95: 715–81.
Williams, Sean Hannon. 2009. Sticky Expectations: Responses to Persistent Over-Optimism in Marriage, Employment Contracts, and Credit Card Use. *Notre Dame Law Review* 84: 733–91.
Wonnacott, Thomas H., and Ronald J. Wonnacott. 1990. *Introductory Statistics*. 5th edn. New York, NY: John Wiley & Sons.
Zamir, Eyal. 1998. The Efficiency of Paternalism. *Virginia Law Review* 84: 229–86.