

Chapter 2

The Value Provided by a Scientific Explanation



Ruobin Gong, Joseph B. Kadane, Mark J. Schervish, Teddy Seidenfeld,
and Rafael Bassi Stern

Abstract In this essay we investigate how deductive and probabilistic explanations add value to the theory providing the explanation. We offer an analysis of how, e.g., an explanation in the form of an abduction differs from a mere prediction in this regard. We apply this analysis to respond to a challenge posed by Glymour (Theory and Evidence. Princeton University Press, Princeton, 1980) regarding Bayesian confirmation of new theories using old data. Last, we consider additional criteria involving subjunctive conditionals and counterfactual conditionals for distinguishing explanations from mere predictions, which help illuminate why an explanation carries different cognitive value than does a mere prediction.

2.1 Introduction

There are varieties of explanatory forms in the sciences. Two forms that are commonly examined are *deductive* and *probabilistic* explanations. (Two classic twentieth century treatises on the philosophy of scientific explanation are Braithwaite's (1953) and Nagel's (1961). Also, see Koslow's (2019, Chap. 8).) In Sects. 2.2 and 2.3 we examine, respectively, how a deductive or a probabilistic explanation adds value to the theory providing the explanation. And we examine how an

R. Gong

Department of Statistics, Rutgers University, New Brunswick, NJ, USA

e-mail: ruobin.gong@rutgers.edu

J. B. Kadane · M. J. Schervish

Department of Statistics, Carnegie Mellon University, Pittsburgh, PA, USA

e-mail: kadane@stat.cmu.edu; mark@stat.cmu.edu

T. Seidenfeld (✉)

Departments of Philosophy and Statistics, Carnegie Mellon University, Pittsburgh, PA, USA

e-mail: teddy@stat.cmu.edu

R. B. Stern

Departamento de Estatística, University of São Paulo, São Paulo, Brazil

e-mail: rbstern@gmail.com

explanation differs from a mere prediction in this regard. In Sect. 2.4 we apply this analysis to respond to a challenge posed by Glymour (1980) regarding Bayesian confirmation of new theories using old data. In Sect. 2.5 we consider additional criteria for distinguishing explanations from mere predictions, which addresses why an explanation carries different cognitive value than does a mere prediction.

2.2 Deductive-Nomological¹ [D-N] Explanations

2.2.1 Three Kinds of D-N Explanations

Nagel's first illustration of the deductive model of explanation is a derivation of an elementary arithmetic generality: A derivation of the generality that the sum of the first k odd integers is the perfect square k^2 . What is being explained here, the *explanandum*, is a specific arithmetic law that applies to each of infinitely many cases. The explanation is a deduction of that law from a set of more general mathematical laws that serve as the premises of a logical argument.²

A more ambitious version of this kind of explanation is found in Sect. 9 of Hilbert's (1971) familiar model of Euclidean plane geometry E using a field of algebraic numbers A : the countable set of numbers arising from finitely many applications of addition, subtraction, multiplication, and division, along with recursion over $|\sqrt{(1+\omega)^2}|$, starting with the integer $\omega = 1$. In Hilbert's model, all of traditional Euclidean plane geometry E is explained (i.e., derived) with this approach, including the parallel postulate. In this model, Euclidean points $\mathbf{A}, \mathbf{B}, \dots$, are identified with ordered pairs of algebraic numbers (x, y) . Lines $\mathbf{a}, \mathbf{b}, \dots$ are identified with ratios of triples of algebraic numbers $(u : v : w)$, where not both $u = 0 = v$. To say that the point (x, y) lies on the line $(u : v : w)$ means that $ux + vy + w = 0$, etc. Then Euclidean plane geometry E , the *explanandum*, is shown to be relatively consistent with a countable model of algebraic numbers A , which entails E under the translation scheme noted above.

A third kind of deductive explanation is illustrated by a species of what Peirce (1955, Chap. 11 in the 1955 collection) calls *abduction*. Here is the deductive form of abduction: A surprising fact F is noted. F cannot be explained based on settled background assumptions, which contributes to F 's status as a *surprise*. Hypothesis H is proposed to explain F , where H is comprised of lawlike generalities that, together with settled background assumptions provide a deductive explanation of F .

¹ We follow common usage that scientific laws are more than mere "accidental" generalizations, and the phrase 'nomological universals' designates the added status. See, e.g., Nagel (1961, Sect. 4.1).

² We imagine that explanation might run as follows. Use the more general arithmetic law that the sum of the first k positive integers, $1 + 2 + \dots + k$, equals $k(k+1)/2$ to show that the sum of the first k positive odd integers, $\sum_{n=0}^{k-1} (2n+1)$, equals $2 \cdot [\sum_{n=1}^{k-1} n] + k = 2[(k-1)k/2] + k = k^2$.

Then H is made worthy of further assessment, e.g. H is now worthy of testing with new, experimental data, because of its value as a potential explanation of F .

What is gained by such deductive explanations? The answer depends upon which question is asked. If the underlying question is, e.g., how to determine the sum of the first k odd integers, the deductive explanation may include a schema for computation. If the underlying question is about consistency of, e.g., Euclidean plane geometry, the deductive explanation may show relative consistency, just as Hilbert showed using a reduction of Euclidean geometry to another mathematical theory whose primitives, algebraic numbers, do not include geometric concepts. And as a bonus, computational methods in the reducing theory, e.g., the determinant of n linear equations in n variables, can be used in higher dimensional geometry to identify the dimension and the volume of the parallelotope determined by those n linear equations: where the parallelotope is the induced mapping of the n -dimensional unit square. And the same deductive explanation can provide the answer to more than one of these questions.

If the question is in the form “Why the surprising fact F ?” then H is a candidate for a deductive explanation and, so, H rises to the status of being a test-worthy hypothesis. H might also allow prediction of a future F -episode, which provides one schema for testing H . It is our purpose in this paper to understand the value of such deductive explanations noting that, in each of these three cases, the value provided by a deductive explanation does not require uncertainty about the explanandum.

2.2.2 *Explanation and Prediction: A Necessary Condition for an Explanation*

Important for our purposes, however, is to distinguish an explanation from a mere prediction (or post-diction). In the case of an explanation, at least one of the premises essential to the derivation of the explanandum is *lawlike* or *nomical*. The contrast is between a generalization, e.g., G : All A 's are B 's, and the enhanced claim that G also is *lawlike*.

We follow Braithwaite's and Nagel's proposals for distinguishing these as follows. A necessary condition N for a generalization G to rise to the status of a lawlike assertion is that

[N] Either G is a postulate of a theory T or is explained within that theory.

That is, as a necessary condition for G to serve as a law in an explanation that, e.g., a specific A is a B , is the requirement, N that G is a consequence within T of some higher level lawlike generalizations, or is fundamental to T . (The postulates of T are assumed lawlike.) Otherwise, if G fails this condition, it provides merely for a prediction that a particular A also is a B . But then G does not *explain* the A - B pattern. Then it merely provides reason to predict that an A is a B .

When considering empirical theories, N is not a sufficient condition for *lawlikeness*, as is illustrated using an example voiced by Russel (1921, Lecture

5) in connection with his concerns about non-uniqueness of causes. (See, also Braithwaite's 1953, pp. 306–8 discussion of this example.) Here is Russell's example, adapted to our purposes.

We seek an explanation for why the workers at a late nineteenth Century Factory #1 go to lunch at approximately noon on workdays. The intended explanation is a derivation of this pattern of behavior from two lawlike generalizations, H and P :

H : Factory #1's horn sounds at about noon on workdays.

and

P : Workers at Factory #1 know it is lunchtime when they hear the factory horn.

Russell's example presumes a commonsense background theory T_1 of Industrial Organization that includes these two generalizations, H and P , as lower level generalizations about Factory #1. (Hence, each of H and P satisfies condition N.) Theory T_1 quantifies over various classes of factories, their methods of communicating, and workers. By design, the factories in a class have a similar organization.

Suppose that this "Russellian" background theory T_1 includes the assumption that Factory #2 is organized for communicating lunch times to its workers just as is Factory #1. The two factories belong to the same class. Their clocks and horns are locally powered and independently coordinated with a common local time. Assume, further, that the background assumptions for the example include the commonplace fact of where the factories are located, and that Factory #2 is located 50 miles to the north of Factory #1. The two factories are located well out of the range of the other's horns. But then, as Russell noted, T_1 also entails the following generalization H' .

H' On workdays, the workers at Factory #1 go to lunch when the horn sounds at Factory #2.

Modified slightly from Russell's original point, the generalization H' also satisfies the condition N for *lawlikeness*, as it too is a deductive consequence of theory T_1 . But unlike the lawlike generalizations H and P , intuitively, H' is *not* lawlike. Though H' is adequate for predicting when workers at Factory #1 go to lunch on a typical workday, it fails to satisfy relevant *subjunctive conditionals* that also are required, we believe, if a generalization is to serve as an explanation. We formulate one such subjunctive conditional as follows:

S If the horn at Factory #2 *were* to sound at time t within the 1/2 hour interval 11:45 AM–12:15 PM, then the workers at Factory #1 *would* go to lunch at time t .

There are at least two relevant ways we understand that background theory T_1 defeats the subjunctive conditional, S . Based on T_1 , we expect that each of the following obtains:

First, if on a workday, by an intervention, the horn at Factory #2 is made to sound at 11:45 AM, then the workers at Factory #1 do not go to lunch at about 11:45, though they may still go to lunch at about noon that day.

Second if, on a workday, unknown to the workers at Factory #1, their horn is deactivated, but the horn at Factory #2 sounds at noon, as usual, then the workers at Factory #1 do not go to lunch at about noon that day. In those circumstances, we expect that workers at Factory #1 experience a delay in going to lunch.

Note how these two cases differ. The first intervention involves directly modifying the event that is a premise of H' (namely, that the horn sounds at Factory #2 at time t) without modifying the event to be explained. Background theory T_1 entails that the conclusion to the generalization H' fails. The second case is to intervene by modifying an event that we believe correctly explains why the workers at Factory #1 go to lunch at about noon (namely, that the horn sounds at Factory #1 at about noon), without modifying the event that serves as the purported explanation in the dubious account H' . That is, in the second case, without intervening on the operations of the horn sounding at Factory #2 at about noon, but by intervening to prevent the horn at Factory #1 from sounding at about noon, we alter the normal lunchtime behavior of the workers at Factory #1.

Of course, these evaluations of what would be the consequences of the imagined interventions depend upon the particulars of the background theory T_1 . For instance, if instead the background theory T_2 includes an assumption that each intervention that makes the horn at one factory sound at a time t also makes the horn at the other factory sound at the same time t , then H' would satisfy the requisite subjunctive conditionals. In that case, we do not see an objection to the assessment that H' provides an explanation, not merely a prediction. We return to these themes in Sect. 2.5.

2.3 Probabilistic Explanations

Next, we discuss explanations that involve probabilities. We follow, e.g., Hempel (1965, Chap. 12.3), in distinguishing *Deductive-Statistical* [D-S] from *Inductive-Statistical* [I-S] explanations.³

D-S explanations are a species of D-N explanations where the *explanandum* E is a statistical law, and the explanation is a derivation of E from observations and some higher-order laws, at least one of which also is statistical. A D-S explanation appeals to the mathematical laws of probability among other laws. For instance, consider the elementary Hardy-Weinberg law for the stable recurrence relation, over

³ Here, we do not address cases of *Direct Inference*: see Levi (1980, Chap. 12). In Direct Inference, the premise of the reasoning is a statement of a chance distribution over outcomes on a kind of trial, and the conclusion is a conditional credence distribution for outcomes on an instance of that kind of trial.

successive generations after the first, of expectations for proportions of alleles and genotypes of a dichotomous Mendelian trait. This law is a deductive consequence (using the laws of conditional probability) of the basic Mendelian laws for a trait that does not mutate, is not sex-linked, and has random mating. (See Hardy (1908).) For our purposes in distinguishing mere predictions from explanations, there are no new issues raised by D-S explanations that are not already evident with D-N explanations. We illustrate this, further, in Sect. 2.5 (below).

I-S explanations, however, involve a different form, with the following canonical schema.

As before, the proposition to be explained is the *explanandum* E , and the premises for the argument are laws L_1, \dots, L_n and observations O_1, \dots, O_m , where at least one of the laws is statistical, i.e. a probability distribution, perhaps indexed with a parameter. To say that these laws and observations explain E is to claim that there is a sound *inductive inference* that concludes E from these laws and observations. And, for our discussion here, further, we distinguish two varieties of inductive inference:

Probabilistic Inductive Inference [PII] The inductive inference in question is formulated as a function solely of the “posterior” conditional probability: $P(E|L_1, \dots, L_n, O_1, \dots, O_m)$. That is, the laws and observations explain E provided that, e.g., the conditional probability $P(E|L_1, \dots, L_n, O_1, \dots, O_m)$ is sufficiently high.

We note that PII inductive explanations admit an evident comparative relation using the criterion of positive/negative relevance. Given $\{L_1, \dots, L_n, O_1, \dots, O_m\}$, say that O_{m+1} is *positively* (respectively, *negatively*) relevant to E provided that

$$P(O_{m+1}|E, L_1, \dots, L_n, O_1, \dots, O_m) > (\text{resp., } <) P(O_{m+1}|L_1, \dots, L_n, O_1, \dots, O_m).$$

Then

$$P(E|L_1, \dots, L_n, O_1, \dots, O_m, O_{m+1}) > (\text{resp., } <) P(E|L_1, \dots, L_n, O_1, \dots, O_m)$$

if and only if

E is positively (respectively, negatively) relevant to O_{m+1} , given $\{L_1, \dots, L_n, O_1, \dots, O_m\}$.

This concept of comparative probabilistic inductive inference is the subject of Glymour’s (1980, Chap. 3) problem of “old evidence,” which we address in Sect. 2.4.

A second kind of I-S explanation uses *decision-theoretic* principles to formulate rules of inference for concluding E from $\{L_1, \dots, L_n, O_1, \dots, O_m\}$. We call these *Decision Theoretic Inductive Inference* [DTII]. For example, in response to the question *Whether E ?* a Bayesian DTII rule uses a cognitive expected utility to rank the three epistemic options: *to accept E* , *to reject E* , or *to suspend judgment about E* . (See Levi (1980, Chap. 2) for the development of one such program.)

What makes the utility cognitive rather than, say, economic, is that the options are formulated as epistemic acts (*what to believe*) rather than in terms of, e.g., the monetary cost of acting on that belief. The outcomes of the epistemic decision are assessed according to scientific goals in coming to believe a proposition, E . They are not assessed by, e.g., the monetary costs for applying E in order to solve a practical problem. These cognitive goals include these familiar considerations: valuing true beliefs over false ones; valuing more informative answers to questions over less informative ones (including the scope of application of an hypothesis as an aspect of informativeness); and valuing the explanatory power of a theory.

Let the cognitive choice be between two rival hypotheses E_1 and E_2 . Allowing that other epistemic considerations are roughly equal between these two, suppose that E_1 explains some significant relevant phenomenon whereas E_2 does not. Then E_1 carries greater expected cognitive utility than E_2 . The preferred option is to conclude E_1 rather than E_2 even if the conditional credences, given $\{L_1, \dots, L_n, O_1, \dots, O_m\}$ are equal: Even if $P(E_1|L_1, \dots, L_n, O_1, \dots, O_m) = P(E_2|L_1, \dots, L_n, O_1, \dots, O_m)$, still E_1 is the preferred conclusion.

2.4 Confirming a New Theory with Old Evidence

Clark Glymour's influential monograph Glymour (1980), *Theory and Evidence*, contains a provocative Chap. 3, *Why I Am Not a Bayesian*, that poses what has become known in the Philosophy literature as the problem of *old evidence*. Glymour proposes the following puzzle: Suppose E stands for evidence known at the time a new theory T is proposed. In the historical example that Glymour uses, we are to reason from the perspective of an informed physicist in 1915. Evidence E is the surprising advance of the perihelion of Mercury—accepted and well known since at least 1865 and surprising relative to Newtonian theory. Its advance was in excess of Newtonian prediction by about 43" per century. T is Einstein's then novel *General Theory of Relativity* [GTR], which in its 1915 published form accounted for the advance of the perihelion.

From within an idealized Bayesian perspective, let probability $P_{1915}(\cdot)$ represent the physicist's unconditional degrees of belief, and let $P_{1915}(\cdot|\cdot)$ be her/his associated conditional probability function. Glymour challenges a PII Bayesian account of theory confirmation as follows. Show that, given Einstein's 1915 explanation, the probability of *GTR* is thereby increased. Glymour assumes the answer is to show that E is positively relevant to *GTR*: $P_{1915}(GTR|E) > P_{1915}(GTR)$. But since E is already known, presumably $P_{1915}(E) = 1$. Then how can E support any change in the probability of the theory *GTR*? A simple calculation shows $P_{1915}(GTR|E) = P_{1915}(GTR \text{ and } E)/P_{1915}(E) = P_{1915}(GTR)$. So, Glymour claims, Einstein's 1915 explanation of E by *GTR* cannot serve as confirmation of *GTR*. This Glymour takes to be a fault in Bayesian analysis.

Here, we do not review the details of the Philosophy literature spawned by this elegant example. But there is a central theme running through a large swath of the

would-be Bayesian replies to Glymour's challenge. We point to Howson (1991) as making that theme clear. Rather than taking $P_{1915}(\cdot)$ as the appropriate credence function for formalizing Einstein's contribution, instead substitute a modified credence, $P_{1915}^*(\cdot)$, which uses a revised corpus of background assumptions that does not include E . So, $P_{1915}^*(E) < 1$. The idea then is that Einstein's contribution is represented as establishing $P_{1915}^*(E|GTR) \approx 1$. If N stands for the rival Newtonian Theory, then also $P_{1915}^*(E|N) < 1$. Assuming, as reasonable, that $0 < P_{1915}^*(GTR), P_{1915}^*(N) < 1$, an elementary calculation establishes the desired confirmation. Then, $P_{1915}^*(GTR|E) > P_{1915}^*(GTR)$.

Glymour (1980, pp. 87–91) anticipated this rebuttal and, we think cogently, argued that an appeal to a “counterfactual” credence, $P_{1915}^*(\cdot)$, revised from an historically accurate $P_{1915}(\cdot)$, does not work. We know of no theory of belief revision that answers the central question: What degrees of belief, circa 1915, should a physicist have held on the counterfactual supposition that deleted what was known, then, for 50 years about the conflict between a Newtonian model of our solar system and the orbit of Mercury around the sun? As the historical record makes plausible, Einstein revised his formulation of *GTR* until it would explain E . (See Earman and Glymour (1978, p. 300).) Then, in Peircean terms, Einstein's 1915 contribution is *abductory* for *GTR*, and not *confirmatory* of *GTR*.

To better understand our response to the *old evidence* problem, perhaps a bit of background on Bayesian analysis is useful at this point. Our Bayesian analysis starts with a person who (at least hypothetically) makes a decision. The decision might be on cognitive issues (*What theory to believe?*) or on financial issues (*What stock to buy?*). Central to Bayesian theory is that decision making follows the criterion of maximizing subjective utility. When the decision maker faces uncertainty about something crucial to the decision, then and only then do probabilities enter the analysis. Thus, in cases where uncertainty is not the driving issue, utilities are far more important to Bayesian analysis than are probabilities. We propose this as the appropriate perspective to adopt when trying to understand the old evidence problem. In brief, our proposal is that Einstein's new 1915 *explanation* of the old evidence E by theory *GTR* increases the utility of *GTR*. Einstein's 1915 contribution does not change a Bayesian analysis of the probabilistic credence (circa 1915) of *GTR*.

In order to apply this line of analysis to Glymour's case, let us suppose that the decision maker in question is a physicist (circa 1915) wanting to understand Mercury's orbit around the Sun. Newtonian theory does not explain the advance, and attempts to save Newtonian theory by postulating additional planetary masses did not succeed.⁴ By contrast, *GTR* offers a much better explanation. But that

⁴ See Roseveare (1982). See Jeffreys (1973, pp. 170–171) and Levenson (2015) for discussions of the historical account about the failure to defend Newton by speculating an unobserved planet, Vulcan, orbiting between Mercury and the Sun. The combination of Newton Theory and the Vulcan hypothesis provided an abductory explanation for the observed advance of Mercury's perihelion. However, that theory also entailed failed predictions, at odds with *GTR*, that were refuted by evidence known by 1915.

explanation is in the form of an abduction, not a prediction. So, Einstein's 1915 explanation contributes to the cognitive value of *GTR*, but not to an increase in its posterior probability—even assuming it is meaningful to assign personal probabilities to whole theories. Indeed, our analysis does not depend on whether or not one tries to answer the question whether it is meaningful to assign probabilities to theories.

In order to appreciate the lawlike status of *GTR* (circa 1915) note that it offered predictions in the form of subjunctives relating to other phenomena, (e.g. the “bending of light”) that were subject to test with the 1919 Dyson-Eddington expedition. Circa 1915, these predictions show that *GTR* had lawlike status. Post 1920, with the addition of the 1919 observations, those predictions added to the utility of *GTR* by establishing its added scope, regardless whether one argues, also, that those observations increased the posterior probability of *GTR*.

Did Einstein's 1915 explanation offer any advice to this physicist on which theory to use in general? Newton's calculations are simpler. But Einstein's might be better, as they are more accurate. If the problem is mere prediction (not explanation) of a mundane event on the surface of the earth—such as pertaining to loads on a traffic bridge—then use Newton. But if you want to know about the effects of a black hole on its surrounding masses, then Einstein's model looks relevant. That Mercury is the closest planet to the Sun is a hint that it is particularly the right kind of situation where Einstein's model will excel. What was to be learned from Einstein's 1915 work is the usefulness of *GTR* as an abductive explanation for some phenomena that elude explanation from the rival, Newtonian Theory.

2.5 Subjunctive Conditionals and Probabilistic Explanations

2.5.1 A Probabilistic Version of Russell's Example

In order for an empirical theory to have explanatory utility, taken together with the settled background assumptions, it must support a derivation of relevant phenomena, must not contradict relevant background knowledge, and also include generalizations that answer relevant subjunctive questions. We discuss these aspects of explanation within a probabilistic version of Russell's example.

In this variant, we consider two rival statistical theories, denoted as T_1^S and T_2^S , which parallel their deductive versions, T_1 and T_2 , of Sect. 2.2.2. Theories T_1^S and T_2^S agree on some (unconditional) probabilistic generalizations: For instance, each of these two theories assigns high probability to the events that the horns at Factories #1 and #2 sound at noon on workdays, and then the workers at the two factories go to lunch. They differ about some conditional predictions. For instance, theory T_1^S models the factory horns as operating nearly independently of each other given the time of day. Theory T_2^S models the horns as conditionally dependent, regardless the time of day.

We formalize the differences between the probabilities assigned by these two theories as follows.⁵ Let H_i be a Bernoulli random variable that indicates whether or not the horn at Factory # i sounds on a particular workday. Let L_i be the time at which the workers at Factory # i go to lunch. And let C_i be the clock time that the horn at Factory # i sounds on the condition that $H_i = 1$. Both theories agree that normally the horns at the factories sound simultaneously at noon and the workers go to lunch. Both theories treat the factories similarly in this way and assign probabilities:

$$P(H_1 = H_2 = 1) = 1 - \epsilon .$$

And for $i = 1, 2$

$$P(C_i = L_i = \text{noon} | H_i = 1) = 1 - \delta$$

where δ and ϵ are small positive quantities. It then follows that the two theories agree on the high probability, conditional prediction: Given the horn sounds at Factory #2, then it is noon and the workers at Factory #1 go to lunch:

$$P(L_1 = \text{noon} | H_2) \geq (1 - \delta)(1 - \epsilon) .$$

Nonetheless, according to theory T_1^S , if the horn at Factory # i fails to sound by noon or very shortly thereafter, regardless whether or when the horn at the other factory sounds, the workers at Factory # i eventually get restless and go to lunch at some time after noon, say, within the interval 12:10 PM–12:20 PM. Theory T_1^S assigns conditional probabilities in accord with the “Russellian” background assumptions about the local configuration of the factories’ horns. It allows the horn at one factory to fail while the other factory’s horn continues to operate. Though there may be scenarios where horns at both factories fail for a common reason, e.g., when there is a region-wide power failure, nonetheless according to T_1^S that joint failure is rare. For $i \neq j$, theory T_1^S proposes:

$$\epsilon < P_1(H_i = 0 | H_j = 0) \leq 2\epsilon .$$

And since workers respond only to the horn they hear,

$$P_1(L_1 \geq 12:10 \text{ PM} | H_1 = 0) = P_1(L_1 \geq 12:10 \text{ PM} | H_1 = 0 \text{ and } H_2 = 1) \approx 1 .$$

⁵ We use subscripts on the probability function to indicate a difference between the two theories. Where the theories assign the same probabilities, we avoid adding a subscript.

Then, theory T_1^S supports the subjunctive conditional prediction that,

Were the horn at Factory #1 to fail ($H_1 = 0$) while the horn at Factory #2 *were* to sound as usual ($H_2 = 1$, $C_2 = \text{noon}$), then the workers at Factory #1 *would* not go to lunch at noon.

The second theory T_2^S , which we view as controversial, also assigns high probability to the events of the horns at both factories sounding simultaneously at noon on workdays, and that the workers go to lunch when the horn at their respective factory sounds. But T_2^S includes conditional probabilities that the horns operate as fully dependent variables: For times, t , $11:45 \text{ AM} \leq t \leq 12:15 \text{ PM}$:

$$P_2(\text{Horn at Factory \#1 sounds at time } t | \text{Horn at Factory \#2 sounds at time } t) \approx 1$$

$$P_2(\text{Horn at Factory \#1 fails to sound at time } t | \\ \text{Horn at Factory \#2 fails to sound at time } t) \approx 1 .$$

Then, the following subjunctive conditional prediction offered by theory T_2^S is not supported by T_1^S .

If 10 minutes before noon the horn at Factory #2 *were* to sound, then workers at Factory #1 *would* go to lunch 10 minutes early, at 11:50 AM.

Advocates of T_2^S might use the conflicts in assessments of these two subjunctive conditionals to cast doubt on the “Russellian” background assumptions: Duhemian underdetermination applies here since we are contemplating what amounts to a crucial experiment for testing T_1^S versus T_2^S .⁶ How to proceed?

According to these theories, as δ and ϵ are very small quantities, it is highly improbable that an investigator will see a situation where $H_1 \neq H_2$, as contemplated in the first of these two subjunctive conditionals, or see a situation where C_2 is earlier than noon, as contemplated in the second. Both theories agree that the situation for testing either subjunctive conditional is a rare event. Then, it is important to examine the opportunity for an experimental *intervention* that creates an otherwise improbable event, but one that is relevant to testing the predictions associated with the subjunctive conditionals. Before giving the details in the factory example, we should be clear about what we mean by observations and interventions.

⁶ See Duhem’s (1916/1953) well known criticism of “crucial experiments.”

2.5.2 Observations and Interventions

A familiar constraint in framing a decision problem for an agent is to distinguish states from options of choice. States are objects of uncertainty for the agent. We distinguish two categories of decision problems:

- (i) decisions where all the agent's options are probabilistically independent of the states, and
- (ii) decisions where states are probabilistically dependent on some of the agent's options.

As we explain below, this distinction is relevant to our view about which subjunctive conditionals *lawlike* generalizations are required to address.

Consider the first category of decisions. As an illustration, suppose the decision is about which observations to make in order to learn a state of Nature: this is an epistemic problem. For example, suppose the decision is about how many siblings to observe in order to learn the genotype of a random F_2 offspring that results from a self-fertilizing hybrid F_0 plant. Then the agent's background (Mendelian) theory stipulates a joint probability distribution over the genotypes of the F_2 generation. Moreover, the agent assesses independence between the choice of how many siblings to observe and the distribution over F_2 genotypes. (Of course, her/his probability for the genotype of the randomly selected F_2 offspring need not be independent of possible observations made on the siblings.)

What is important for our discussion here is that when there is probabilistic independence between states and choices, then also the agent has a coherent unconditional joint probability distribution on observations of random variables defined with those states.⁷ Let X and Y be random variables whose values are determined by the states. Let the decision problem be of the first kind (i), i.e., where the agent assesses states and choices as independent. So, the agent holds an unconditional joint distribution P over X and Y . Then she/he may address subjunctive conditionals of the form: If random variable X were to satisfy $X = x$, then what *would be* the uncertainty about Y ? The answer is $P(Y|X = x)$. And if the agent's decision is whether to observe variable X , because of the independence of states and choices, then also $P(Y|X = x)$ is the agent's opinion about Y were she/he to choose to observe X and learn that $X = x$. This is our account of how to assess subjunctive conditionals with decisions of the first category: what we call the case of observations. Lawlike generalities are required to support subjunctive conditionals with decisions that involve observation variables.

By definition, the second category of decisions (ii) involve some extent of "act/state dependence," including what economists call cases of "moral hazard." Then, the agent's uncertainty about states of Nature may depend upon which option she/he chooses. If as in a controlled experiment, the agent has the option to intervene

⁷ In Savage's (1954) theory, all decision problems are of this first kind. In his decision theory, states and options are probabilistically independent.

and directly to fix the value $X = x$, then her/his uncertainty about Y given the choice to fix $X = x$ may well differ from her/his uncertainty about Y when, also, there is the opportunity to choose merely to observe X and then to learn that $X = x$. In cases of intervention, we use the conditional probability distribution, $P(Y | \text{fix } X = x)$ to assess the subjunctive conditional: What *would be* the uncertainty about Y *were* the agent to intervene and to fix $X = x$?

In contrast with the foregoing analysis of subjunctive conditionals for category (i) decisions, in category (ii) decisions we do not grant coherence of an unconditional joint probability over random variables defined by states of Nature. Our reasons for this difference relate to the admittedly controversial issue about whether to allow the agent probabilistic uncertainty about her/his current choices.⁸ We find compelling Savage's position that while contemplating a decision problem, the agent's uncertainty about what to choose is *not* representable through the values of her/his personal probability function, which guides decision making. Thus, when deciding on an experiment design, we do not recognize as meaningful for the experimenter an unconditional probability for, e.g., $P(\text{fix } X = x)$. Hence, when contemplating a category (ii) decision about whether or not to fix $X = x$, the agent has no unconditional probability for the state $X = x$. This is because the event " $X = x$ " is coextensive with the union of the two disjoint events $\{\text{fix}(X = x), \text{do not fix}(X = x)\}$. Of course, our position here is consistent with the agent having well defined personal probabilities about some future decisions she/he might face, and more controversially and given the fallibility of memory, also about some past decisions that he/she made.

This prohibition against unconditional probability distributions over states in category (ii) decisions means that a subjunctive conditional of the form: "What would be the uncertainty for Y given that $X = x$ is fixed?" is assessed directly by the parallel conditional probability distribution, $P(Y | \text{fix } X = x)$, which values are not determined by the conditional probability distribution $P(Y | X = x)$ that arises from unconditional probabilities that are well defined in category (i) decisions. Then, in order to defend the lawlike status of a generality in category (ii) problems, specifically in order to provide analysis of relevant subjunctive conditionals when interventions are introduced as the conditioning event, the agent's credences require more conditional probabilities than are determined by unconditional credences.

⁸ For some background on this controversy, see Levi (2007). For a contrary position to the one we endorse here, see Jeffrey (1965). In Jeffrey's theory, there is no corresponding difference in the decision theoretic assessments for what here are called "options" and "states." In Jeffrey's theory each proposition is assigned both a probability and a utility.

2.5.3 Interventions in the Factory Example

Theory T_1^S incorporates relevant predictions in experiments that introduce *interventions* (“experimental treatments”) on the operations of factory horns. The *interventions* are of two kinds: One kind of intervention compels a factory horn to sound at a designated time. The other kind of intervention prevents the horn from sounding at a designated time. According to the “Russellian” background, and as supported by T_1^S if, for instance, an intervention silences the horn at noon at Factory #i, then the workers at that factory react to the absence of the horn and go to lunch somewhat later than normal that day. That is, given an intervention where the horn at Factory #i is silenced at noon, then T_1^S stipulates the following workers’ behaviors.

$$P(H_i = 0 | \text{intervene at noon to prevent the horn sounding at Factory \#i}) = 1$$

and

$$P(L_i > \text{noon} | \text{intervene at noon to prevent the horn sounding at Factory \#i}) = 1.$$

So, if one contemplates an intervention in which the horn at Factory #1 fails to sound at noon, while there is no intervention made on the horn at Factory #2, we achieve a test case for one of the relevant subjunctive conditionals that is necessary for determining whether it is T_1^S or T_2^S that explains the workers’ ordinary lunchtime behaviors. Likewise, T_1^S contemplates an intervention that compels the horn at Factory #2 to sound earlier than noon, while leaving unchanged the operation of the horn at Factory #1. This affords a test of the second subjunctive conditional prediction on which the two theories disagree.

What this example illustrates is an important methodological result: Rival theories may also raise controversies about what counts as the settled background assumptions for an inquiry. The possibility of experimental interventions permits empirical testing of rival predictions of some relevant subjunctive conditionals, independent of what may be these controversial background assumptions. That is, in order to assess the explanatory content of the two rivals, T_1^S and T_2^S , experimenters may be able to sidestep a concomitant disagreement about what to admit as the background assumption regarding, e.g., the organization of the factories.

If, in accord with the “Russellian” background assumptions, the experimenter can intervene to prevent the horn at Factory #1 from sounding at noon, while the horn at Factory #2 operates normally, then we have devised a test of the T_1^S prediction that the workers at Factory #1 have a delayed lunch that day, while the workers at Factory #2 go to lunch as usual, at noon. And if that prediction accords with the workers’ behavior, then T_1^S satisfies the requirement of supporting at least this subjunctive conditional.

On the other hand, it might turn out that the intervention fails. The experimenter might be unable to implement the experimental design of intended “treatments”

to the two horns. It might be that each attempt to silence the horn at Factory #1 also silences the horn at Factory #2. Each attempt to sound the horn early at one factory also sounds the horn simultaneously at the other factory. Then the advocates of T_2^S gain momentum for their proposal to revise the “Russellian” background assumptions about how the two horns are controlled. Then theory T_1^S does not succeed in supporting a subjunctive prediction that we take to be necessary for explaining, rather than merely predicting, the workers’ lunchtime behavior.

We summarize: In general, a probabilistic theory needs to specify the joint distributions for a set of relevant variables. Such joint distributions are specified conditional on all of the background information assumed by the theory. The relevant variables, whose joint distributions are to be specified, include both passively *observed variables* and *intervention variables* that are set deliberately—if any exist. When interventions are feasible, they permit empirical testing of some subjunctive conditional predictions that (we propose) are necessary in order for a theory to provide scientific explanations, and thereby gaining added cognitive utility over rival theories that merely predict (passive) observables.

2.5.4 *Equilibria and Subjunctives*

No doubt, the reader will recognize there is a simple *causal structure* in the elementary probabilistic versions T_1^S of Russell’s example of the two factories, as reviewed in Sect. 2.5.1: The sounding of a factory horn shortly before noon is a direct cause of those factory workers recognizing that it is lunchtime. Under normal conditions, these causes have a common cause, e.g. the time of day. However, given the common cause, the direct causes are independent. Hence, given the state of the horn at Factory #1, the operation of the horn at Factory #2 is independent of the behavior of the workers at Factory #1. That causal structure is made evident through, e.g., the directed statistical graphs used in the elegant, important contemporary theories of causal inference created by Pearl (2009) and Spirtes et al. (2001). However, scientific laws do not presume a causal structure. We illustrate with the Hardy-Weinberg Law, an instance of a D-S explanation.

The Hardy-Weinberg Law provides necessary and sufficient conditions for stable recurrence of expectations for genotypes through successive generations. Its form is an equilibrium in expectations for genotype proportions over successive generations, not as a causal mechanism for achieving those genotype proportions.

For simplicity, we deal with an elementary case of a dichotomous Mendelian trait. Each organism in the population carries two genes for a given trait, each of which may be either the dominant allele form **A**, or the recessive form **a**. Thus, there are three genotypes: the homozygous dominant form **AA**, the hybrid form **Aa**, and the homozygous recessive form **aa**. Under basic Mendelian theory, the dominant form **AA** and the hybrid form (**Aa**) are phenotypically indistinguishable. However, that assumption is not required for deriving the Hardy-Weinberg Law, below.

Suppose at generation F_0 , the population proportions of these three genotypes types are, respectively, $\langle x_0, y_0, z_0 \rangle$. Let \mathbf{p}_0 be the proportion of dominant genes in the F_0 population, and \mathbf{q}_0 the proportion of recessive genes. That is, $\mathbf{p}_0 = (2x_0 + y_0)/2$. And so $\mathbf{q}_0 = (y_0 + 2z_0)/2$. Assume random mating and the absence of selectivity for the three phenotypes.⁹ This includes assuming that there is no sex-linkage for the trait, and no selective advantage among the genotypes for the trait. Then by deductive reasoning from these elementary probabilistic assumptions, it follows that the expected proportions of genotypes in the F_1 generation is given by $\langle x_1, y_1, z_1 \rangle$, satisfying $x_1 = p_0^2 = x_0^2 + x_0y_0 + (y_0/2)^2$, $z_1 = q_0^2 = z_0^2 + z_0y_0 + (y_0/2)^2$, and $y_1 = 2\mathbf{p}_0\mathbf{q}_0 = y_0(1 - y_0/2) + 2x_0z_0$. Thus, the expected proportion of dominant genes in the F_1 population, \mathbf{p}_1 , satisfies $\mathbf{p}_1 = (2x_1 + y_1)/2 = (2\mathbf{p}_0^2 + 2\mathbf{p}_0\mathbf{q}_0)/2 = \mathbf{p}_0(\mathbf{p}_0 + \mathbf{q}_0) = \mathbf{p}_0$, which is stationary from the previous generation. And, by the Law of Conditional Expectations,¹⁰ this recurs for the expectation of gene proportions in the F_2 and subsequent generations.

But this result does not establish recurrence of the expectations of the three genotype proportions between the F_0 and F_1 generations. The Hardy-Weinberg Law identifies necessary and sufficient conditions for recurrence of the expectations for three genotypic proportions. The Hardy-Weinberg condition (2.1) for this recurrence between one generation, say the j^{th} generation, F_j , and its subsequent generation F_{j+1} , is that the parental generation, F_j , satisfies:

$$y_j^2 = 4x_jz_j. \quad (2.1)$$


This equality may fail for the F_0 generation, which then is not in Hardy-Weinberg equilibrium. But under the stated assumptions (random mating, etc.) it obtains for the expectations in the F_1 generation where $y_1^2 = 4\mathbf{p}_0^2\mathbf{q}_0^2 = 4x_1z_1$. Hence, under the stated assumptions, the expectations relating the F_1 and F_2 generations, and (again by the Law of Conditional Expectations) also the expectations relating genotypic proportions in subsequent generations after F_1 , satisfy the Hardy-Weinberg condition (2.1).

Next, we use a de Finetti diagram, Fig. 2.1 (below) to represent the Hardy-Weinberg Law. (See Edwards (2000) for a general discussion of this technique.) A population is individuated by its proportions for the three genotypes. The state-space for these genotype proportions is represented by the probability simplex over three states, with extreme points, $\langle 1, 0, 0 \rangle$, $\langle 0, 1, 0 \rangle$ and $\langle 0, 0, 1 \rangle$. Each population proportion is represented by a unique point in this state space. Condition (2.1) for

⁹ The assumption of random mating is not satisfied, for instance, in Mendel's classic experiments with pea plants. Pea plants are self-fertilizing—pollen is not randomly scattered—which feature was essential for Mendel's experimental design. Then, over successive generations, the two homozygous types are absorbing and the hybrid type is transient.

¹⁰ Specifically, by the *Law of Total Probability*, expectations and conditional expectations for bounded random variables X and Y satisfy, $E[X] = E[E[X|Y]]$.

Fig. 2.1 Hardy-Weinberg equilibrium in a de Finetti diagram



recurrence of the expectations of genotype proportions in successive generations is depicted in Fig. 2.1 by the convex (quadratic) curve of genotypic ratios.


A “prior” distribution over the points on this curve corresponds to a (de Finetti) exchangeable distribution over sequences of genotypes from a hypothetically infinite sequence of draws, one from each generation. This probabilistic model provides a basis for answering some subjunctive conditional questions, in accord with the constraint for *lawlikeness* that we require.

Q: Were the population proportion in generation F_j , \mathbf{p}_j for the dominant allele **A** to drift to a new value, for instance, suppose that $\mathbf{p}_j = 1/4 \neq \mathbf{p}_0 = 1/2$, what *would be* the expected genotype proportions in subsequent generations after the j^{th} ?

The populations that satisfy a subjunctive condition of the form $\mathbf{p}_j = x_j + y_j/2 = \mathbf{c}$ (for $0 \leq \mathbf{c} \leq 1$) are represented in the de Finetti diagram by a vertical line with lower endpoint on the bottom face of the simplex, $\langle \mathbf{c}, 0, (1 - \mathbf{c}) \rangle$, and upper endpoint either $\langle 0, 2\mathbf{c}, (1 - 2\mathbf{c}) \rangle$, if $\mathbf{c} \leq 1/2$, or $\langle 2\mathbf{c} - 1, 2(1 - \mathbf{c}), 0 \rangle$, if $\mathbf{c} \geq 1/2$, which is on one or the other of the upper faces of the simplex. By the Hardy-Weinberg Law we derive that, for a population that satisfies the new condition, $\mathbf{p}_j = \mathbf{c}$ and conditions of random mating, etc., then beginning with generation F_{j+1} the expected proportions for the three genotypes are $\langle \mathbf{c}^2, 2\mathbf{c}(1 - \mathbf{c}), (1 - \mathbf{c})^2 \rangle$. Note this makes the process Markovian in that the answer to the expected proportions in generations subsequent to the j^{th} does not depend upon the population’s prior historical state. The value of \mathbf{p}_0 is irrelevant. This, then, answers the subjunctive conditional question using an ordinary derivation from the Hardy-Weinberg Law.

Figure 2.2, below, graphs the Hardy-Weinberg solution, for $\mathbf{c} = \frac{1}{4}$, as the intersection of the vertical line ($\mathbf{p} = \frac{1}{4}$), which depicts all the populations that satisfy the antecedent of the subjunctive conditional, and the Hardy-Weinberg Law (2.1). That point of intersection is $\langle 1/16, 3/8, 9/16 \rangle$.

Fig. 2.2 Graph of populations with gene proportion $p = \frac{1}{4}$



Though the Hardy-Weinberg Law fixes expectations for the three proportions of genotypes in subsequent generations once the new condition $p = c$ is satisfied, it does not provide the causal probabilistic dynamics (as might be given with an acyclic directed graph) for how this equilibrium is achieved. Each population that satisfies the new condition, $p = \frac{1}{4}$ in the example of Fig. 2.2, achieves the same equilibrium in expectations for proportions of genotypes in subsequent generations. In Fig. 2.2, each population that falls on a given vertical line achieves the same equilibrium after one generation.

We do not claim that the Hardy-Weinberg Law is inconsistent with a “causal” theory. One might also provide the causal statistical dynamics for how the population proportions move in time through the state-space. The conclusion we argue for here is that causal structure is not necessary for a law to be explanatory.

2.5.5 Counterfactual Conditionals

In augmenting Braithewaite’s and Nagel’s condition **N** for *lawlikeness* with an additional requirement, that in order for a generalization to explain and not merely to predict an outcome, we require the generalization to support reasoning with *subjunctive conditionals*—either relating to hypothetical *interventions* or to possible *observations*. However, we do not require that for a generalization to be a *law* also that it supports *counterfactual* conditional reasoning.

To make the point explicit with Russell’s example of the two factories, we require that intervening to make the horn at Factory #2 sound at 11:45 is appropriate for testing the subjunctive offered by theory T_2^S ,

“If the horn at Factory #2 *were* to sound at 11:45 AM, the workers at Factory #1 *would* then go to lunch earlier than usual.”

Since we expect this subjunctive to fail, we expect the problematic generalization to fail the necessary condition for lawlikeness. Namely, we expect that though the generalization

“The workers at Factory #1 go to lunch when the horn at Factory #2 sounds”

may be useful for prediction, it does not support the relevant subjunctive conditional; hence, we expect it does not pass muster as an explanation, since it fails being lawlike.

But we have not required as test for *lawlikeness* that a generalization also supports *counterfactual* conditional reasoning. We do not require support for conditionals of the form: “Yesterday, *had* the horn at Factory #2 sounded at 11:45 then the workers at Factory #1 *would have* gone to lunch early.” By stipulation, the antecedent to this conditional is known to be false: It is accepted that, yesterday, the horn at Factory #2 sounded at noon, as usual. The antecedent for this conditional is inconsistent with the decision maker’s background knowledge. So, the conditional probabilities of Sect. 2.5.2 do not apply, since conditional probability is not defined relative to an inconsistent set of assumptions.

In order to apply those conditional probabilities to the specific event in question, the decision maker needs to revise her/his background assumptions. She/he has to remove (to “contract”) the assumption that the horn at Factory #2 sounded at noon yesterday, in order to make logical space for a hypothetical event. But the intelligibility of that maneuver is precisely what Glymour questions in connection with the attempt to show that (circa 1915) Einstein’s new explanation for old data conforms to a case of Bayesian credal confirmation of *GTR*. And we agree with Glymour’s skepticism about the intelligibility of such a revision in the settled assumptions about Mercury’s orbit.

There may be specialized circumstances where it is evident how to carry out a revision of settled background knowledge so as to “test” a counterfactual conditional. For instance, there may be no controversy accepting the counterfactual that, though in fact there was no drift in the gene frequency between the F_0 and F_1 generations, say $\frac{1}{2} = \mathbf{p}_0 = \mathbf{p}_1$, *had there been* drift to $\mathbf{p}_1 = \frac{1}{4}$, then the expectations for the F_2 genotypes *would have been* at the values $< 1/16, 3/8, 9/16 >$. The answer in this case can rest on the assumption of exchangeability in the statistical model. Then the counterfactual question about a past event is treated as an ordinary subjunctive question about the outcome of a future event whose value is not fixed by background assumptions. It may be that in legal settings counterfactual reasoning of this kind is required in order to resolve issues of liability about specific grievances. (See, e.g., Stern and Kadane (2019).)

However, we are skeptical of a general theory of counterfactual conditionals. For instance, suppose that there is no genetic drift in the F_1 generation—say $1/2 = \mathbf{p}_0 = \mathbf{p}_1$ —but there *is* observed drift in the F_2 generation, with $\mathbf{p}_2 = 1/4$. How to evaluate the counterfactual conditional that, *had there been* drift in F_1 to $\mathbf{p}_1 = 1/3$, then there *would have been* a second instance of drift in F_2 , with $\mathbf{p}_2 = 1/4$, matching the F_2

value that was observed.¹¹ Our view is that there does not yet exist a functioning account of counterfactual reasoning that rises to the level needed for making testing counterfactual conditionals into a well-formed criterion of *lawlikeness*.

2.6 Summary

In this essay we explore the idea that a scientific theory T may gain in cognitive value for us when we learn that T *explains* a specific event E , even though E is already known. One of our concerns is to understand how an explanation differs in value from a mere prediction. For that purpose, we explore probabilistic explanations that support various forms of reasoning with subjunctive conditionals: involving antecedents to conditionals that use either interventions on, or observations of random variables. We allow that the event E that is explained by T may already be known but, as in cases of Peircean abduction, E may be surprising because we do not understand it. In such a case there is no puzzle how old data may play a role in increasing the value of a novel theory. This contrasts with the puzzle of “old evidence” posed by Glymour where, instead, the challenge is to show that the novel explanation of old evidence by a new theory increases the probability or “credence” for that new theory.

References

- Alchourrón, C., Gärdenfors, P. and D. Makinson. 1985. On the logic of theory change: partial meet functions for contractions and revisions. *Journal of Symbolic Logic* 50:510–530.
- Braithwaite, R.B. 1953. *Scientific explanation*. London: Cambridge University Press.
- Duhem, P. 1916/1953. Physical theory and experiment. In *Readings in the philosophy of science* (translated and reprinted), eds. Feigl, H., and M. Brodbeck, 235–252. New York: Appleton-Century-Crofts.
- Earman, J., and C. Glymour. 1978. Einstein and Hilbert: two months in the history of general relativity. *Archive for History of Exact Sciences* 19:291–308.
- Edwards, A.W.F. (2000) *Foundations of mathematical genetics*, 2nd edn. Cambridge: Cambridge University Press.
- Glymour, C. 1980. *Theory and evidence*. Princeton: Princeton University Press.
- Hardy, G.H. 1908. Mendelian proportions in a mixed population. *Science* 78:49–50.
- Hempel, C.G. 1965. *Aspects of scientific explanation*. New York: The Free Press.
- Hilbert, D. 1971. *Foundations of geometry*, trans of the 10th German edn. Open Court: La Salle, Ill.
- Howson, C. 1991. The ‘old evidence’ problem. *The British Journal for the Philosophy of Science* 42:547–555.
- Jeffrey, R. 1965. *The logic of decision*. New York: McGraw Hill.

¹¹ This example illustrates the ongoing dispute about the so-called “Recovery” postulate in the AGM theory (Alchourrón et al. 1985) of belief revision. See Levi (2004) for helpful discussion.

- Jeffreys, H. 1973. *Scientific inference* (3rd edn). Cambridge: Cambridge University Press.
- Koslow, A. 2019. *Laws and explanations; theories and modal possibilities* (Synthese Library), vol 430. Switzerland: Springer Nature.
- Levenson, T. 2015. *The Hunt for Vulcan . . . and how Albert Einstein destroyed a planet, discovered relativity, and deciphered the Universe*. New York: Random House.
- Levi, I. 1980. *The enterprise of knowledge*. Cambridge: MIT Press.
- Levi, I. 2004. *Mild contraction: evaluating loss of information due to loss of belief*. Oxford: Oxford University Press.
- Levi, I. 2007. Deliberation does crowd out prediction. In *Hommage à Wlodek; 60 Philosophical Papers Dedicated to Wlodek Rabinowicz*, eds. T. Rønnow-Rasmussen, B. Petersson, J. Josefsson, and Egonsson, D. www.fil.lu.se/HomageaWlodek. Last accessed 25 Feb 2022.
- Nagel, E. 1961. *The structure of science*. New York: Hardcourt, Brace, and World.
- Pearl, J. (2009) Causal inference in statistics: an overview. *Statistics Surveys* 3:96–146.
- Peirce, C.S. 1955. *Philosophical writings of Peirce*, ed. J. Buchler. New York: Dover.
- Roseveare, N.T. 1982. *Mercury's perihelion, from Le Verrier to Einstein*. Oxford: Oxford University Press.
- Russel, B.R. 1921. *The analysis of mind*. London: Allen & Unwin.
- Savage, L.J. 1954. *The foundations of statistics*. New York: John Wiley & Sons.
- Spirtes, P., Glymour, C., and R. Scheines. 2001. *Causation, prediction, and search*, 2nd edn. Cambridge: MIT Press.
- Stern, R.B., and J.B. Kadane. 2019. Indemnity for a lost chance. *Law, Probability and Risk* 18:115–148.