

A Little Survey of Induction

John D. Norton¹

Department of History and
Philosophy of Science
University of Pittsburgh
jdnorton@pitt.edu

Prepared for Conference on Scientific Evidence, Center for History and
Philosophy of Science, Johns Hopkins University, April 11-13, 2003;
to appear in P. Achinstein, ed., *Scientific Evidence: Philosophical and Historical
Perspectives* (provisional title).

My purpose in this chapter is to survey some of the principal approaches to inductive inference in the philosophy of science literature. My first concern will be the general principles that underlie the many accounts of induction in this literature. When these accounts are considered in isolation, as is more commonly the case, it is easy to overlook that virtually all accounts depend on one of very few basic principles and that the proliferation of accounts can be understood as efforts to ameliorate the weaknesses of those few principles. In the earlier sections, I will lay out three inductive principles and the families of accounts of induction they engender. In later sections I will review standard problems in the philosophical literature that have supported some pessimism about induction and suggest that their import has been greatly overrated. In the final sections I will return to the proliferation of accounts of induction that frustrates efforts at a final codification. I will suggest that this proliferation appears troublesome only as long as we expect inductive inference to be subsumed under a single formal theory. If we adopt a material theory of induction in which individual inductions are licensed by particular facts that prevail only in local domains, then the proliferation is expected and not problematic.

1. Basic Notions

Inductive inference is the primary means through which evidence is shown to support the content of science. It arises whenever we note that evidence lends support to an hypothesis, in whatever degree, while not establishing it with deductive certainty. Examples abound. We note that some squirrels have bushy tails and infer that all do. We note that our cosmos is bathed in a 3K radiation bath and this lends credence to the standard big bang cosmology that predicts it. Or we may infer to the limited

¹ I am grateful to Peter Achinstein for helpful discussion and comments on an earlier draft and to Phil Catton for helpful discussion.

efficacy of a drug if the health of members of a test group given the drug improves more than that of a control group denied the drug. This notion of induction is called "ampliative," which means that the hypotheses supported are more than mere reformulations of the content of the evidence. In the simplest case, evidence that pertains to a small number of individuals is generalized to all. It is called "ampliative," since we amplify a small number to all. This modern tradition differs from an older tradition that can be traced back to Aristotle in which induction meant generalization and the generalization need not necessarily be ampliative. In that tradition, taking a finite list (Iron conducts electricity; gold conducts electricity;...) and summarizing it (All metallic elements conduct electricity.) is called a "perfect induction," even though it is fully deductive.

In cases in which we find the evidence so compelling that we accept the hypothesis into our corpus of belief, we tend to use terms "induction" or "inductive inference." If the bearing of the evidence is weak and we merely want to register that it has lent some support to the hypothesis, then we more commonly use the term "confirmation." When the evidence speaks against an hypothesis but does not disprove it, we count it as a case of "disconfirmation." The introduction of the term "confirmation" into the inductive inference literature is relatively recent and closely associated with the representation of inductive relations as probabilistic relations. If the degree of confirmation rises to a level at which we are willing to adopt the hypothesis without proviso, then we say that the hypothesis is detached from the evidence; if it is wanted, a formal theory of confirmation requires a "rule of detachment" to implement it. For detachment in a probabilistic account of induction, the rule would typically require the probability to rise above a nominated threshold.

2. Three Principles and the Families They Engender

Any account of inductive inference specifies when an item of evidence inductively supports an hypothesis and, in some cases, provides a measure of the degree of support. There are so many such accounts, some of them with histories extending to antiquity, that it is impossible to discuss them all. The task of comprehending them is greatly simplified, however, once we recognize that virtually all accounts of induction are based on just three ideas. As a result, it is possible to group virtually all accounts of induction into three families. This system is summarized in Table 1 below. Each family is governed by a principle upon which every account in each family depends. I also list what I call the "archetype" of each family. This is the first use of the principle and a familiar account of induction its own right. These archetypes suffer weaknesses and the family of accounts grows through embellishments intended to ameliorate these weaknesses.

<i>Family</i>	Inductive Generalization	Hypothetical Induction	Probabilistic Induction
<i>Principle</i>	An instance confirms the generalization.	The ability to entail the evidence is a mark of truth.	Degrees of belief are governed by a numerical calculus.
<i>Archetype</i>	Enumerative induction	Saving the phenomena in astronomy.	Probabilistic analysis of games of chance
<i>Weakness</i>	Limited reach of evidence	Indiscriminate confirmation	Applicable to non-stochastic systems

Table 1. Three Families of Accounts of Inductive Inference

Each of the three families will be discussed in turn in the three sections to follow and the entries in this table explicated.² While most accounts of inductive inference fit into one of these three families, some span across two families. Achinstein's (2001) theory of evidence, for example, draws on ideas from both hypothetical induction and probabilistic induction, in so far as it invokes both explanatory power and probabilistic notions. Demonstrative induction, listed here under inductive generalization, can also be thought of as an extension of hypothetical induction.

3. Inductive Generalization

The Archetype: Enumerative Induction

The most ancient form of induction, the archetype of this family, is "enumerative induction" or "induction by simple enumeration." It licenses an inference from "Some As are B." to "All As are B." Examples are readily found in Aristotle.³ Traditionally, enumerative induction has been synonymous with induction and it was a staple of older logic texts to proceed from deductive syllogistic logic to inductive logics based on the notion of enumerative induction. These include variant forms of enumerative induction such as "example" (This A is B; therefore that A is B.) and "analogy" (a has P and Q; b has P; therefore b has Q.) One variant form is quite subtle. Known as intuitive induction it requires that induction must be accompanied by a "felt certainty on the part of the thinker" (Johnson, 1922, p. 192). Elusive as this notion is, it may have tacitly been part of the notion of enumerative induction as far back as Aristotle. It has been just as traditional to vilify enumerative induction. Francis Bacon (1620, First Book, §105) has the most celebrated jibe:

The induction which proceeds by simple enumeration is puerile, leads to uncertain conclusions, and is exposed to danger from one contradictory instance, deciding generally from too small a number of facts, and those only the most obvious.

² It is a hopeless to try to include every view on induction in a survey of this size. I apologize to those who find their views slighted or omitted.

³ For example, *Prior Analytics* Book II.23 68b15-20 (McKeon, 1941, p.102); *Topics* Book I, Ch. 12 105a10-19 (McKeon, 1941, p. 198).

Part of that scornful tradition has been the display of counterexamples devised to make the scheme appear as foolish as possible. My own view is that the scorn is misplaced and that the counter examples illustrate only that any induction always involves inductive risk. Since the conclusion is never guaranteed at the level of deductive certainty, failures are always possible. The actual inductive practice of science has always used enumerative induction and this is not likely to change. For example we believe all electrons have a charge of -1.6×10^{-19} Coulombs simply because all electrons measured so far carry this charge.

Extensions

The principal weakness of enumerative induction is its very limited scope. It licenses inference just from some A's are B to all and that scheme is too impoverished for most applications in science. Most of the embellishments of the scheme seek to extend the inductive reach of the evidence. In doing this they extract the governing principle from the archetype of enumerative induction: an instance confirms the generalization.

There are two avenues for expansion. The first reflects the fact that enumerative induction is restricted by the limited expressive power of the syllogistic logic in which it is formulated. The principal burden of Hempel's (1945) "satisfaction" criterion of confirmation is to extend this notion of confirmation from the context of syllogistic logic to that of the richer first order predicate logic. Hempel's theory in all detail is quite complicated, but its core idea is simple. Take an hypothesis of universal scope, for example $(x)(Px \rightarrow Qx)$ (i.e. "For all x, if x has P then x has Q.") The development of the hypothesis for a class of individuals is just what the hypothesis would say if these individuals were all that there were. The development of $(x)(Px \rightarrow Qx)$ for the class {a,b} would be $(Pa \rightarrow Qa) \& (Pb \rightarrow Qb)$. Hempel builds his theory around the idea that the development is the formal notion of instance. Thus his account is based on a formal rule that asserts that hypotheses are confirmed by their developments.

Elegant as Hempel's account proved to be, it was still restricted by the inability of instances to confirm hypotheses that used a different vocabulary. So the observation of bright spots in telescopes or on cathode ray tube screens could not confirm hypotheses about planets or electrons. The second avenue of expansion seeks to remedy this restriction. One of the most important is a tradition of eliminative methods most fully developed in the work of Mill in his *System of Logic* (1872, Book III, Ch. 7). The methods are intended to aid us in finding causes. I may find, for example, that my skin burns whenever I give it long exposure to sunlight and that it does not burn when I do not. Mill's "Joint Method of Agreement and Disagreement" then licenses an inference to sunlight as the cause of the burn. There are two steps in the inference. The first is a straightforward inductive generalization: from the instances reported, we infer that my skins always burns just when given long exposure to sunlight. The second is the introduction of new vocabulary: we are licensed to infer to the sunlight as the cause. So the methods allow us to introduce a causal vocabulary not present in the original evidence statements.

Mill's methods extend the reach of evidence since we use theoretical results to interpret the evidence. In Mill's case it is the proposition that a cause is an invariable antecedent. Glymour's (1980) "bootstrap" account allows any theoretical hypothesis to be used in inferring the inductive import of

evidence. For example, we observe certain lines in a spectrograph of light from the sun. We use known theory to interpret that as light emitted by energized Helium; so we infer that our star (the sun) contains Helium, which is an instance of the hypothesis that all stars contain Helium. What is adventurous in Glymour's theory is that the hypotheses of the very theory under inductive investigation can be used in the interpretation of the evidence. Glymour added clauses intended to prevent the obvious danger of harmful circularity. However the ensuing criticism of his theory focussed heavily on that threat.

In Glymour's bootstrap, we use theoretical results to assist us in inferring from the evidence to an *instance* of the hypothesis to be confirmed. In the limiting case, we infer directly to the hypothesis and simply do away with the inductive step that takes us from the instance to hypothesis itself. In the resulting "demonstrative induction" or "eliminative induction" or "Newtonian deduction from the phenomena," we find that the auxiliary results we invoke are sufficiently strong to support a fully deductive inference from the evidence to the hypothesis. While this scheme may have little interest for those developing theories of inductive inference, it has very great practical importance in actual science, since it offers one of the strongest ways to establish an hypothesis. It has been used heavily throughout the history of science. Newton used it repeatedly in his *Principia*; it has also been used often in the development of quantum theory (See Norton, 1993, 2000). It should be stressed that demonstrative induction can never do away with the need for other inductive inference schemes. Since it is fully deductive, it never really allows us to infer beyond what we already know. Its importance lies in alerting us to cases in which the inductive risk needed to move from evidence to hypothesis has already been taken elsewhere. If we have already accept the inductive risk taken in believing the auxiliary results in some other part of our science, we find that we need take no further inductive risks in inferring from evidence to hypothesis.

4. Hypothetical Induction

The Archetype: Saving the Phenomena

This family of accounts of induction is based on a quite different principle: the ability of an hypothesis to entail deductively the evidence is a mark of its truth. The origins of explicit consideration of this principle lie in astronomy. Plato is reputed to have asked his students to find what combinations of perfect circular motions would save the astronomical phenomena. (See, for example, Duhem, 1969.) Whether celestial motions that save the phenomena were thereby shown to be the true motions became a very serious issue when Copernicus showed in the 16th century that the astronomical motions could be saved by the motion of the earth. What resulted was a fierce debate, dragging into the 17th century, over whether these motions were nonetheless merely mathematical fictions or true motions. (See for example, Jardine, 1984.)

What became clear then was that *merely* saving the phenomena was not enough. One could always concoct some odd hypothesis able to save the phenomena without thereby having a warrant to it. We can now present the problem in a compact and forceful way through the example of frivolous conjunction. Assume that some hypothesis H is able to entail deductively (i.e. "save") the evidence E,

usually with assistance from some auxiliary hypotheses. Then it is a simple matter of logic that a logically stronger hypothesis $H' = H \& X$ has the same ability, even though the hypothesis X conjoined to H can be the most silly irrelevance you can imagine. If saving the phenomena is all that matters, then we are licensed to infer from E to H and just as much to H' even though E now indirectly supports the silly X .

A principle had been extracted from the archetype of saving the phenomena in astronomy: the ability of an hypothesis to entail the evidence is a mark of its truth. But the unaugmented principle accorded the mark too indiscriminately. So a simple hypothetico-deductive account of induction—that is one that merely requires the hypothesis to save the evidence⁴—is not a viable account of induction. This one problem, directly or indirectly, drives the literature in this family. It seeks to embellish the simple account with some additional requirement that would tame its indiscriminateness.

Extensions

The most straightforward embellishment produces what I shall call "exclusionary accounts." In them, we require that the hypothesis H entail the evidence E and moreover that there is some assurance that E would not have obtained had H been false. Thus competitors to H are excluded. In the simplest version we merely require that the evidence E in conjunction with suitable auxiliaries entails deductively the hypothesis. This is immediately recognizable as the "demonstrative induction" introduced above, although now the alternative term "eliminative induction" is more appropriate because we recognize the power of the inference to eliminate alternative hypotheses. While this kind of deductive exclusion is less commonly possible, one can often show something a little weaker: that, were H false, then *most probably* E would not have obtained. This circumstance arises routinely in the context of controlled studies. Randomization of subjects over test and control group are designed to assure us that any systematic difference between test and control group (the evidence E) must be due to the treatment (the hypothesis H), since, if H were false, the differences could only arise by vastly improbably coincidences. This model of traditional error statistical analysis drives such accounts of induction as Giere (1983) and the more thorough account of Mayo (1996).

Other exclusionary accounts draw on our quite vivid intuitions concerning quite vague counterfactual possibilities. Nearly a century ago, Perrin found that roughly a dozen independent experimental methods for determining Avogadro's number N all gave the same result. In what Salmon (1984, Ch. 8) calls a "common cause argument," Perrin argued that this is powerful evidence for the reality of atoms, for—and here is the counterfactual supposition—were atoms not real, it would be highly unlikely that all the experimental methods would yield the same value.⁵ Related accounts of induction

⁴ Here and henceforth I will tacitly assume that the hypothesis entails the evidence usually with the aid of auxiliary hypotheses.

⁵ This same argument form, specifically involving multiple experiments to measure some fundamental constant, has arisen in other areas, so that I have given it the narrower label of the method of overdetermination of constants. (Norton, 2000, Section 3)

have been developed under the rubric of "common origin inferences" (Janssen, 2002) and Whewell's "consilience of induction."

A different approach attempts to tame the indiscriminateness of hypothetico-deductive confirmation by using the notion of simplicity. Of the many hypotheses that save the phenomena, we are licensed to infer to the simplest. (See for example, Foster and Martin, 1966, part III.) This may seem a fanciful approach given the difficulty of finding principled ways to discern the most simple. However in practice it is much used. The most familiar usage comes in curve fitting. While many curves may be adequate to the data points, allowing for experimental error, we routinely infer to the simplest. What makes a curve simpler is usually quite precisely specified; in the family of polynomials, the simpler are those of lower order. The preference for simplicity is so strong that standard algorithms in curve fitting will forgo a more complicated curve that fits better in favor of a simpler curve that does not fit the data as well.

In an account of inductive inference known as "abduction" or "inference to the best explanation" (Harman, 1965; Lipton, 1991) we tame the indiscriminateness of simple hypothetico-deductive inference by requiring the hypothesis not just to entail the hypothesis but to explain it. We infer to the hypothesis that explains it best. The 3K cosmic background radiation is not just entailed by big bang cosmology, it is also quite elegantly explained by it. We would prefer it to another cosmology that can only recover the background radiation by artful contrivance that we do not find explanatory. Just as choosing the simplest hypothesis threatened to enmesh us in a tangled metaphysics of simplicity, this account brings us the need to explicate the notion of explanation. There is a quite expansive literature already on the nature of explanation; in the context of abduction, causal explanation—explaining by displaying causes—seems favored. In practical applications, identifying the simplest explanation can often be done intuitively and without controversy. One gets a sense of the dangers lurking if one considers a putatively successful experiment in telepathy. A parapsychologist would find the experimental result best explained by the truth of telepathy. A hard boiled skeptic (like me) would find it best explained by some unnoticed error in the experimental protocols.

In what I shall call "reliabilist accounts," merely knowing that an hypothesis saves the evidence is insufficient to warrant support for it from the evidence; in addition we have to take into account how the hypothesis was produced. It must be produced in the right way; that is, it must be produced by a method known to be reliable. One of the best known reliabilist accounts is incorporated in Lakatos' (1970) "methodology of scientific research programs." According to it, theories arise in research programs and continued pursuit of the program is warranted by the fecundity of the program. A program scoring successful novel predictions is "progressive" and worthy of pursuit; whereas one without is languishing—"degenerating"—and might be abandoned. Decisions on theory evaluation must be made in the context of this history. Merely noting a static relationship between the theory and a body of evidence is not enough. This structure is already in place in Popper's (1959) celebrated account of scientific investigation proceeding through a cycle of conjecture and refutation. The newly conjectured hypothesis that survives serious attempts at falsification is "well corroborated," a status that can only be assigned in the light of the history of the hypothesis' generation. As part of his complete denial of inductive inference,

Popper insisted that the notion of corroboration was quite different from confirmation. I have been unable to see sufficient of a difference to warrant the distinct terminology.

Reliabilist approaches permeate other assessments of the import of evidence. We routinely accept the diagnostic judgements of an expert, even when we laypeople cannot replicate the expert's judgements from the same evidence. We do this since we believe that the expert arrived at the assessment by a reliable method, perhaps learned through years of experience. Reliabilism also underwrites our scorn for *ad hoc* hypotheses. These are hypotheses that are explicitly cooked up to fit the evidence perfectly. For example, in response to failed experiments to detect any motion of the earth in the light carrying ether, we might hypothesize that the earth just happens to be momentarily at rest in it. We do not doubt that the hypothesis entails the evidence; but we doubt that the evidence gives warrant for belief in the hypothesis exactly because of the history of the generation of the hypothesis.⁶

5. Probabilistic Induction

The Archetype: Games of Chance

Accounts in this third family owe their origins to an advance in mathematics: the development of the theory of probability, starting in the 17th century, as a means of analyzing games of chance. It was recognized fairly quickly that these probabilities behaved like degrees of belief so that the same calculus could be used to govern degrees of belief in inductive inference. The best way to represent the inductive import of evidence E on hypothesis H was to form the conditional probability $P(H | E)$, the probability of the hypothesis H given that E is true. One of the most useful results in the ensuing inductive logic was presented by Bayes:

$$P(H | E) = (P(E | H) / P(E)) P(H)$$

It allows us to assess how learning E affects our belief in H, for it relates $P(H)$, the prior probability of H, to $P(H | E)$ the posterior probability of H conditioned on E. These quantities represent our degree of belief in H before and after we learn E. To compute the change, Bayes' formula tells us, we need only know two other probabilities. $P(E)$ is the prior probability of E, our prior belief in the evidence E obtaining whether we know H is true or not. The likelihood $P(E | H)$ is often readily at hand. Loosely it tells us how likely E would be if H were true; in case H entails E, it is one. For more details of this Bayesian approach to inductive logic, see Howson and Urbach (1993) and Earman (1992).

This probabilistic approach is at its strongest when we deal with stochastic systems, that is, those in which probabilities arise through the physical properties of the systems under consideration. These physical probabilities are sometimes called "chances." In such cases we might well be excused for failing to distinguish our degree of belief from the chance of, say, the deal of a royal flush in a game of poker or the decay of a radioactive atom sometime during the time period of its half life. The weakness lies in the

⁶ Kelly (1996) has developed a general framework for reliabilist accounts using the theoretical apparatus of formal learning theory. The framework extends well beyond the simple notion here of reliabilism as an augmentation of hypothetico-deductivism.

difficulty of knowing how to proceed when our beliefs pertain to systems not produced by stochastic processes. Why should simple ignorance be measured by degrees that conform to a calculus devised for games of chance? Why should I even expect all my many ignorances to admit the sort of well behaved degrees that can always be represented by real numbers between 0 and 1?

Extensions

The archetype of the family is the probabilistic analysis of games of chance. While it is not clear that exactly this archetype can be applied universally, a family of accounts grew based on the notion that this archetype captured something important. It is based on the principle that belief comes in degrees, usually numerical, and is governed by a calculus modeled more or less closely on the probability calculus. The weakness addressed by the different accounts of the family is that degrees of belief are not the chances for which the calculus was originally devised.

There have been two broad responses to this weakness. The first is the majority view among those in philosophy of science who use probabilities to represent beliefs. They urge that, while chances may not be degrees of belief, the latter should nonetheless always be governed by the same calculus as chances. In support of this view, they have developed a series of impressive arguments based loosely on the notion that, if our beliefs are to conform to our preferences and choices, then those beliefs must conform to the probability calculus. The best known of these arguments are the “Dutch book” arguments of de Finetti. In a circumstance in which we place bets on various outcomes, they urge that, were our degrees of belief not to conform to the probability calculus, we could be induced to accept a combinations of bets that would assure us of a loss—a Dutch book.⁷ (See de Finetti, 1937; Savage, 1972.)

This majority view has traditionally been accompanied by zealous efforts to give a precise interpretation of the notion of probability employed. Many candidates emerged and could be grouped into roughly three types: a physical interpretation based on relative frequencies of outcomes; a logical interpretation representing probabilities as degrees of entailment; and a subjective interpretation representing probabilities as conventionally chosen numbers constrained by the probability calculus. (For a recent survey, see Gillies, 2000.) While an appeal for precision in meaning is important, in my view, these demands became excessive and placed impossible burdens on the analysis that cannot typically be met by accounts of central terms in other theories. A demand that probabilistic belief be defined fully by behaviors in, for example, the context of betting is reminiscent of long abandoned efforts to define all

⁷ The weakness of all these arguments is that one cannot recover something as highly structured as the probability calculus from an argument without including assumptions of comparable logical strength in its premises. These assumptions are introduced through presumptions about the context of our decisions, the structures of our preferences and the rules for translating beliefs into actions. A skeptic recognizes that these assumptions are essentially generated by working backwards from a foregone conclusion—that beliefs must in the end conform to the probability calculus—to a set of assumptions carefully contrived to encode just that calculus. The skeptic wonders why the entire exercise could not be abbreviated simply by positing that beliefs are probabilities at the start.

physical concepts in terms of measuring operations or psychological states in terms of behaviors. Similarly, efforts to find a non-circular definition for probability in terms of relative frequencies seem to neglect the obvious worry that all terms cannot be given explicit, non-circular definitions without reintroducing circularity into the total system.⁸ One can only guess the disaster that would have ensued had we made similarly stringent demands for precise definitions of terms like the state vector of quantum theory or energy in physics.

The second broad response to the weakness of the family is to accept that degrees of belief will not always conform to the probability calculus and that weaker or even alternative calculi need to be developed. Take for example some proposition A and its negation $\neg A$ and imagine that we know so little about them that we just have no idea which is correct. On the basis of symmetry, one might assign a probability of $1/2$ to each. But that, roughly speaking, says that we expect A to obtain one in two times and that seems to be much more than we really know. Assigning say 0.1 or even 0 to both seems more appropriate. But that is precluded by the condition in probability theory that the probabilities of A and $\neg A$ sum to one. The simplest solution is to represent our belief state not by a single probability measure but by a convex set of them.⁹ More adventurously, one might start to construct non-additive theories, such as the Shafer-Dempster theory (Shafer, 1976). Once one starts on this path, many possibilities open, not all of them happy. Zadeh's "possibility theory" for example computes the possibility of the conjunction (A and B) by simply taking the minimum of the possibilities of each conjunct. The loss of information is compensated by the great ease of the computational rule in comparison with the probability calculus. (For a critique, see Cheeseman, 1986.) My contribution to this literature is the "theory of random propositions." It is a strict weakening of the probability calculus with well defined semantics designed to show that alternative calculi can solve some of the notorious problems of the Bayesian system. The theory of random proposition, for example, is not prone to the problem of the priors; assigning a zero prior probability no longer forces all of its posterior probabilities to be zero.

6. Properties and Tendencies

Once we identify these three families, it is possible to see some properties and tendencies peculiar to each family. One of the most important is that inference in the family of inductive generalization tends to proceed from the evidence to the hypothesis; we start with "Some A's are B" and infer to "All A's are B." In the family of hypothetical induction, the direction of inference is inverted; we first show that one can pass deductively from the hypothesis to the evidence and then one

⁸ If an outcome has probability $1/2$, then in the limit of infinitely many repetitions the frequency of success will not invariably approach $1/2$; rather it will approach $1/2$ *with probability 1*—so a definition based on this fact becomes circular.

⁹ For example we might take $P_1(A)=0.1$ $P_1(\neg A)=0.9$ and $P_2(A)=0.9$ $P_2(\neg A)=0.1$ and choose as our set all probability measures that can be formed as suitably weighted averages of them: $P_\alpha = \alpha P_1 + (1-\alpha)P_2$, where $0 \leq \alpha \leq 1$.

affirms that the evidence supports the hypothesis. Therefore, assigning evidence to the depths and theory to the heights, I label inductive generalization as “bottom up” and hypothetical induction as “top down.”

This difference inclines the two families to very different attitudes to induction. In inductive generalization, the distance between evidence and theory is small so there is optimism about the power of evidence. The passage from evidence to theory is almost mechanical. We replace a “some” by an “all”; or we methodically collect instances and apply Mill’s methods to reveal the cause. This family is hospitable to logics of discovery, that is, to recipes designed to enable us to extract theoretical import from evidence. In hypothetical induction, the distance between evidence and theory is great. There is no obvious path suggested by the evidence to the hypothesis that it supports. We cannot pass from the 3K cosmic background radiation to big bang cosmology by a simple rule. This invites a focus on a creative element in confirmation theory: we can only confirm hypotheses if we are independently creative enough to find the ones able to entail the evidence in the right way. Thus this family is the more traditional home of pessimism over the reach of evidence and the underdetermination thesis to be discussed below. It also invites skepticism about the possibility of logics of discovery.

There is also an interesting difference between, on the one hand, inductive generalization and hypothetical induction and, on the other, probabilistic induction. In the former, there tends to be little sophistication in justifying the particular schemes. They are either taken to be self evident, much as modus ponens in deductive logic is not usually justified. Or they are made plausible by displaying case studies in which the scheme at issue is seen to give intuitively correct results. The justification of schemes is a great deal more elaborate and sophisticated in probabilistic induction, with serious attempts to support the chosen system that go beyond mere invocation of intuitive plausibility, such as the Dutch book arguments sketched above.

7. The Problems of Induction

Stated

Science is an inductive enterprise and its unmatched success must be counted as a triumph of inductive inference. Yet philosophers have traditionally found it hard to participate in the celebration. The reason is that inductive inference has been the customary target of skeptical philosophical critiques. Here I sketch three of the best known. I will also indicate why I think none of them present difficulties anywhere near as severe as is commonly assumed.

Hume’s problem of induction a.k.a. *The problem of induction*. (Salmon, 1967, p.11) This problem, whose origin traces at least to Hume, identifies a difficulty in justifying any inductive inference. It asserts that any justification of a given inductive inference necessarily fails. If the justification is deductive in character it violates the inductive character of induction. If it uses inductive inference, then it is either circular, employing the very same inference form in the justification, or it employs a different inductive inference form, that in turn requires justification, thereby triggering an infinite regress. So we cannot say that our next inductive generalization will likely succeed since our past inductive generalizations have

succeeded, for that is an induction on our past successes that uses exactly the inductive scheme under investigation.

Grue (Goodman, 1983) In this problem, an application of instance confirmation purports to show that the observation of green emeralds can confirm that as yet unobserved emeralds are blue. The trick is to define the predicate “grue,” which applies to emeralds if they are observed to be green prior to some future time *t* and blue otherwise. That some emeralds were observed to be green is equally correctly described as the observation of grue emeralds. The observation confirms all emeralds are green; in its redescribed form, it confirms all emeralds are grue, that is, that as yet unobserved emeralds are blue. One might be inclined to block the confirmation of “all emeralds are grue” by complaining that “grue” is a bogus, compound property not amenable to inductive confirmation. Goodman argues, however, that this complaint fails. He takes grue and an analogously defined “bleen” and uses them as his primitives properties. Green and blue are defined in terms of them by formula just like those used for grue and bleen, so they now appear equally bogus. There is a perfect symmetry in their mutual definitions.

The underdetermination thesis. (Newton-Smith, 2001) In the case of cutting edge science, we are used to several theories competing, with none of them decisively preferred, simply because sufficient evidence has not yet been amassed. The underdetermination thesis asserts that this competition is ineliminable even for mature sciences. It asserts that any body of evidence, no matter how extensive, always fails to determine theory. The thesis is grounded largely in the remark that many distinct theories can save any given body of evidence and in the possibility of displaying distinct theories that share the same observational consequences.

Answered

Let me briefly indicate why none of these problems is so severe. While Hume’s problem has generated an enormous literature, what has never been established is that inductive inference is the sort of thing that needs to be justified. If the very notion of a mode of justification is to be viable, we cannot demand that all modes must in turn be justified, on pain of circularity or infinite regress. Some modes must stand without further justification and inductive inference, as a fundamental mode of inquiry seems as good a candidate as any. While I believe that settles the matter, some may find this response to be question begging.¹⁰ In discussion below, I will suggest another less expected resolution.

In practice gemologists have not been surprised by the continuing green color of emeralds. The reason is that they recognize, perhaps tacitly, that green is a natural kind property that supports induction whereas grue is not.¹¹ What of the symmetry of grue/bleen and green/blue? Our judgment that green only is a natural kind term breaks the symmetry. We might restore it by extending the “grue-ification” of all predicates until we have a grue-ified total science. We now have two sciences with green a natural kind term in one and grue a natural kind term in the other. However I have urged elsewhere

¹⁰ This is one of several standard answers to the problem. For an entrance into this quite enormous literature, see Earman and Salmon, 1992, §§2.5-2.6.

¹¹ Again this is one standard response drawn from an enormous literature. See Stalker (1994).

that, with this extension, we now have two formally equivalent sciences and cannot rule out that they are merely variant descriptions of the very same facts. Then the differences between green and grue would become an illusion of different formulations of the same facts. See Norton (manuscript b).

If the underdetermination thesis is intended to pertain to the reach of evidence by inductive inference, its continued popularity remains a puzzle. For it can only be viable as long one ignores the literature in induction and confirmation. It seems to be based essentially on the notion that many theories can save the same phenomena and thus that they are equally confirmed by it. But, as we saw in some detail in the discussion of hypothetical induction, that notion amounts to the simple hypothetico-deductive account of confirmation that has essentially never been admitted as a viable account exactly because of its indiscriminateness. The underdetermination thesis presumes this indiscriminateness is the assured end of all analysis and neglects the centuries of elaboration of accounts of hypothetical induction that followed, not to mention accounts in the other two families. I have also argued that the possibility of observationally equivalent theories proves to be a self defeating justification for the thesis. If the equivalence is sufficiently straightforward to be established in philosophical discourse, then we can no longer preclude the possibility that the two theories are merely notational variants of the same theory. In that case, the underdetermination ceases to pertain to factual content. (See Norton, manuscript a.)

8. The Nature of Induction

The Problem of Proliferation

While I have reviewed some of the major problems of induction above, there is another neglected problem implicit in the survey. While we think of induction as a sort of logic, it is quite unlike deductive logic in at least one important aspect. Deductive logic proceeds from a stable base of universal schemas. There is little controversy over modus ponens: If A then B; A; therefore B. Any grammatically correct substitution for A and for B yields a valid deduction. In contrast, in the literature on induction and confirmation there is a proliferation of accounts and there seems to be no hope of convergence. Even our best accounts seem to be in need of elaboration. We are happy to infer to the best explanation. But delimiting precisely which inferences that licenses must await further clarification of the nature of explanation and the means for determining what explains best. The Bayesian system works well when our beliefs pertain to stochastic processes; but we are less sure we have the right calculus when the objects of belief stray far from them. Applying standard schemes even in simple cases can reveal gaps. We measure the melting point of a few samples of the element bismuth as 271°C. On their strength, by enumerative induction, we feel quite secure in inferring that all samples of bismuth have this melting point. But had our observation been of the melting point of a few samples of wax or the colors of a few birds of a new species, we would be very reluctant to infer to the generalization. In short, our accounts of induction are proliferating and, the more closely we look at each, the more we are inclined to alter or add further conditions. Induction seems to have no firm foundation; or, if it has, we seem not to have found it.

Has this problem already been solved by one of the accounts outlined above? If there is one account of induction that does aspire to be the universal and final account, it is the Bayesian account.

While its scope is great and its ability to replicate almost any imaginable inductive stratagem is impressive, let me briefly indicate why it cannot yet claim this universal status. There are two of problems. First, probabilities are not well adapted to all situations. For example, they are additive measures and we know from measure theory that some sets are unmeasurable. We can contrive circumstances in which we seek a degree of belief in an unmeasurable set. (See Norton, manuscript, Section 3.)

Second, while it is true that Bayesianism has an impressive record of replication of existing inductive stratagems, this record comes at a cost. We end up positing a fairly large amount of additional, hidden structure somehow associated with our cognition: our beliefs must behave like real numbers or sets of them; we must harbor large numbers of likelihoods and prior probabilities; and we must somehow tacitly be combining them in just the way the probability calculus demands, even though carrying through the calculation explicitly with pencil and paper might well be taxing even for an accomplished algebraist. We might not demand that these extra structures be present tacitly in our actual thought processes. We might merely expect a formal analysis using these structures to vindicate our inductive stratagems. There would still be a difficulty in so far as these extra structures can extend too far beyond the structures of our actual stratagems. This is reminiscent of the situation in geometry a century ago when non-Euclidean geometries began to break into physics. Whatever the geometrical facts, one could always restore Euclidean geometry by adding further hidden geometrical structure. A three dimensional spherical geometry, such as Einstein introduced with his 1917 cosmology, could be constructed by taking a hyperspherical surface in a four dimensional Euclidean space. That a Euclidean geometry could accommodate this new geometry was not decisive; it could only do it at the cost of adding additional structure (an extra spatial dimension) that was deemed physically superfluous. Analogously we need to assess whether our gains from the more adventurous applications of Bayesianism are outweighed by the cost of the additional structures supposed. A Bayesian analysis serves two masters. The structures it posits must conform both to our good inductive stratagems and to the probability calculus. Those dual demands are often answered by the positing of quite rich structures. We saw above that belief states of total ignorance were represented as convex sets of infinitely many probability measures in order not to violate the additivity of the probability calculus. If we drop one of these two masters—that the demand that the structures posited must in the end conform to the probability calculus—then we are more likely to be able to develop an account that posits sparser structures closer to those present in our actual inductive practices.

A Material Theory of Induction

I have argued in Norton (forthcoming) that the problem of proliferation admits a simple solution. The problem arises because we have been misled by the model of deductive inference and we have sought to build our accounts of induction in its image. That is, we have sought formal theories of induction based on universal inference schemas. These are templates that can be applied universally. We generate a valid inductive inference merely by filling the slots of a schema with terms drawn from the case at hand. The key elements are that the schemas are universal—they can be applied anywhere—and that they supply the ultimate warrant for any inductive inference.

If there is such formal theory that fully captures our inductions, we have not yet found it. The reason, I propose, is that our inductions are not susceptible to full systematization by a formal theory. Rather our inductions conform to what I call a "material theory of induction." In such a theory, the warrant for an induction is not ultimately supplied by a universal template; the warrant is supplied by a matter of fact.¹² We are licensed to infer from the melting point of a few samples of bismuth to all samples because of a fact: elements are generally uniform in these physical properties. That fact licenses the inference, but without making the inference a deduction. The inductive character is preserved by the qualification "generally." Contrast this with the case of waxes. We are not licensed to infer from the melting point of a few samples of wax to all samples because there is no corresponding, licensing fact. In general, waxes are different mixtures of hydrocarbons with no presumption of uniformity of physical properties. The general claim is that, in all cases, inductions are licensed by facts. I have given facts that bear this function the name "material postulate."

The idea that the license for induction may derive from facts about the world is certainly not new. Perhaps its best known form is Mill's (1872, Book III, Ch. III) postulation of the "axiom of the uniformity of the course of nature." The difficulty with seeking any one such universal fact to license all inductions is that the fact must be made very vague if it is to pretend to have universal scope, so much so that it is not longer usable. As a result I urge that inductions are always licensed by material postulates that prevail only in local domains; specific facts about elements license inductions in chemistry; specific facts about quantum processes license inductions about radioactive decay (see below); and so on. As a result, I urge that "all induction is local."

Identifying the Material Postulates

The examples of the melting points of bismuth and wax illustrate the principal assertions of the material theory of induction. What grounds do we have for believing that the theory holds for all inductive inference forms and not just the examples presented? Those grounds lie in a review of a representative sample of inductive inference forms to be presented here. We will see that each form depends upon one or other factual material postulate. The ease with which they are identifiable and the way we will see them arising will make it quite credible that material postulates underlie all viable inductive inference forms.

The survey of Sections 2-5 makes it easy for us to sample inductive inference forms from across the literature and affirm that they are all grounded in some sort of local material postulate. (For further discussion, see also Norton, forthcoming, Section 3.) The family of forms encompassed by inductive generalization is the easiest to analyze. All the ampliative inductive inference forms in that family depend upon the same inductive move: an instance confirms the generalization. They differ only in the details of

¹² This is not such a strange notion even in deductive inference. An inference from Socrates' humanity to his mortality is warranted by the fact that all men are mortal. What is different in this example is that (unlike a material theory of induction) the warrant for the inference is in turn supplied by a schema, *modus ponens*.

how the generalization is specified, with Hempel's account the most elaborated. All agree that that an individual that has property P confirms that all individuals have property P, where no restriction is placed on the selection of property P. As familiar problems such as grue have shown, that is far too permissive a condition. Further constraints must be placed on which properties P support generalization. These are routinely factual constraints, such as we saw in the case of bismuth and wax, and these added constraints also typically supply us information of great importance on the strength of the induction.

The basic principle of the family of forms embraced by hypothetical induction is that an hypothesis is confirmed by its deductive consequences. So the hypothesis A&B is confirmed by its consequence A. If the principle is to be non-trivial, that A confirms A&B must amount to more than the trivial A confirms A; it must mean that A confirms B as well. That is only admissible if there is some sort of relevance relation between A and B—for example that that truth of A makes B's truth more likely.¹³ But there is no general theory of what it is for A to be relevant to B; it is introduced as a particular fact essential to the viability of the induction. This pattern continues as we review the various augmented forms that comprise the family. In the exclusionary accounts, we have an added assurance that the evidence would probably not have obtained had the hypothesis been false. That assurance is supplied by factual conditions. For example they are in the factual conditions describing a controlled experiment that assure us that the only systematic difference between test and control group is the application of the treatment. Or they reside in the speculative but widely held belief that, were matter not atomic, Perrin would not have recovered the agreements found in his multiple measurements of Avogadro's number. In another approach, we augment the simple hypothetico-deductive schema by the requirement that the hypotheses in question be simple. There is no universal, formal account of what it is to be simple. Judgments of simplicity, in practice, turn out to reflect our beliefs on the basic ontology of the domain in question. In curve fitting, the natural default is to judge a linear relation simpler than a polynomial. That reflects factual presumptions and these presumptions are made apparent when our simplicity rankings are reversed merely by a change of domain. If we are fitting trajectories to newly observed celestial objects, we will be inclined to fit a conic section ahead of any other curve, even a straight line. Factually, the conic sections are the simplest curves in that new domain; according to Newton's theory of gravity, they are the trajectories of bodies in free fall.

This pattern persists in the remaining inductive inference forms in the family of hypothetical induction. In inference to the best explanation, we augment the hypothetico-deductive schema by the requirement that the hypothesis best explain the evidence. Judgments of what explains and what best explains are governed by factual matters. For example, take the imaginary experiment mentioned above that purports to vindicate telepathy. Judgments of what explains best is essentially determined by the facts we presume to prevail. The parapsychologist, inclined to believe in paranormal processes, explains the outcome of the experiment by the supposition of telepathic processes; the skeptic, disinclined to believe in them, explains it by an unnoticed procedural error. Finally, reliabilist accounts of induction add

¹³ This is a variation of the problem of frivolous conjunction, which is blocked by the factual assertion of relevance.

the requirement that the hypothesis that saves the evidence must be generated by a reliable method. The viability of reliabilist accounts depends essentially on an important factual assumption: that the world is such that methods used do actually lead to the truth. If we do not think that fact obtains, then no inductive gain arises from following the method of hypothesis generation touted as reliable. For this reason, most of us do not heed predictions of oracles and entrail readers, no matter how perfectly they follow their standard methods.

We already saw above that probabilistic accounts of induction work best when our beliefs pertain to stochastic systems. For then we can adopt the physical chances as degrees of belief that automatically conform to the probability calculus. The facts that license the resulting inductions are just the facts that govern the stochastic processes. For example, the half life of Radium-221 is 30 seconds, so that the chance of a decay of a given atom in 30 seconds is 0.5. That fact licenses the corresponding probabilistic degree of belief and, using the calculus to which probabilities conform, we infer the atom will almost certainly decay over the next 300 seconds; the probability of decay is $1-(0.5)^{10}=0.999$. Matters are far less clear when our beliefs pertain to circumstances for which physical chances are unavailable. Then why should our beliefs be captured by numerical degrees and, if they can be, why should those numerical beliefs conform to the probability calculus? A literature mentioned in Section 5 above seeks to argue that our beliefs in this extended domain should still be probabilities. These arguments proceed from factual premises, so it is a simple matter to convert each of these arguments to serve the material theory of induction. These factual premises supply the material postulate for a domain in which a probabilistic account of induction is appropriate. For example, the Dutch book arguments tell us that probabilistic induction is appropriate if we find ourselves in a situation in which we are prepared to accept bets according to the presumptions laid out. Where such arguments are available, we can identify facts that comprise the material postulate; where such arguments are not available, we have no good reason to expect probabilistic induction to be applicable.

More Generally

Seeing this number of examples gives a sense of inevitability for the possibility of the identification of a material postulate for any inductive inference form. The identification can almost be reduced to a general recipe. For each inductive inference form, one merely needs to conceive a factual circumstance in which the form would fail. The presumption that this factual circumstance does not obtain is then declared to be part of the material postulate that licenses the induction.

We can strengthen our expectation that a material postulate can always be recovered by considering how they were recovered in the cases above. In each, the postulate remedies an incompleteness in the specification of the inductive inference form. There are two sorts of incompleteness. Some are incomplete in so far as they simply need factual presumptions for their applicability; there is no expectation that we will supplant these factual presumptions with a general theory. For example, the exclusionary accounts always require some facts to sustain the requirement that the evidence would be unlikely to obtain if the hypothesis were false. Or reliabilist accounts presume that the relevant methods are properly adapted to our world, which is a matter of fact about the methods and the world.

Probabilistic accounts require that we be in a domain in which beliefs can properly be required to conform to the probability calculus. In these cases, straightforwardly, the inductive inference forms can only be employed when the pertinent facts obtain.

In the other sort of incompleteness, the full specification of the inductive inference form calls up general notions from other parts of philosophy of science. Might these full specifications be universal in scope and thus provide counter examples to the material theory of induction's claim that all induction is local? The envisaged counter examples fail. Forming them requires general accounts of these extra notions drawn from other parts of philosophy of science. But these general accounts have, in each case, proven elusive. Instead of seeking them, we cite facts to complete demonstration of the applicability of the inductive inference form to the case at hand. For example, in the case of inductive generalization, we need some way to select which properties support generalization and to which strength. That might come from the notion of natural kind or from notions of physical laws in which those properties may appear as fundamental terms. In the absence of a compact account of natural kinds or physical laws, we cite facts to license the induction: physical properties of elements support inductive generalization; grue-ified properties do not. For simple hypothetico-deductive confirmation, we require a general notion of relevance applicable to parts of a theory. In the absence of a general account of the notion of relevance, we may merely cite the fact that the parts of a theory are relevantly related. When the scheme is augmented by the requirement that the hypotheses to be confirmed must be simple, we do not resort to the ever elusive, general theory of what it is to be simple. We simply declare what is simpler than what and, as we saw above, those declarations are governed by the facts prevailing in the case at hand. Finally in the case of inference to the best explanation, we lack an appropriate account of explanation that could be used to give a full specification of the inductive inference form. If explaining is merely subsumption under a covering law, then the form reverts to simple hypothetico-deductive confirmation. If explanation is the displaying of causes, then we face the problem that there is no generally agreed upon account of the nature of causation that would allow us to specify just what counts as a cause. In practice, however, we do not need a general account of explanation. In each case, we believe certain facts prevail and provide guides to the sorts of results that may govern the process at hand. We typically pick from this repertoire the hypotheses that we judge to explain the evidence and which do it best. In short, in all these cases we may speak of general notions like law, natural kind, simplicity and explanation; but in practice we ground the induction not in the general theory of these notions but in particular facts.

For these reasons, the material theory of induction urges that all inductive inference is ultimately licensed by facts prevailing only in local domains. Indeed the theory presents a challenge to critics: can we find a universally valid inductive inference schema? The material theory predicts that this is not possible.

The Problems It Solves

This locality of induction resolves the problem of proliferation. Facts in each domain license inductions unique to that domain. There are many cases in which inductions licensed in different domains appear similar. The inference from the melting point of some samples of bismuth to all samples is formally like the inference from the blackness of some ravens to all. We codify this with a schema:

enumerative induction. But the similarity is superficial since the two inductions are ultimately licensed by very different facts. The first induction is quite strong and second somewhat weaker, since color uniformities in birds are far less reliably obtained than the uniformities of physical properties of elements. This difference of strength is a puzzle for the scheme of enumerative induction. It is expected according to a material theory since the inductions are from different domains where they are licensed by different facts. As we restrict the domain further, the particular facts available to license induction become more specific and varied. As a result we expect this sort of closer scrutiny to produce a greater variety of formally different sorts of inductions with no presumption of a convergence to a few identifiable forms. This divergent behavior was identified as the problem of proliferation of formal theories; from the perspective of a material theory, the proliferation is an unproblematic artifact of the nature of induction.¹⁴

Finally a material theory of induction casts a very different light on Hume's problem of induction. That problem, as we saw above, arises from our ability rapidly to generate difficulties in seeking a justification of an induction. Our justification is circular, in so far as we use the same inductive inference form for justification. Or it triggers an infinite regress, in so far as we must always invoke another inductive inference form to justify the last one. The regress is problematic even in its early stages. It is fanciful to carry out a meta-induction on inductions; and still more fanciful to try to carry out meta-meta-inductions on meta-inductions on inductions—and that has taken us only two steps into the regress. A material theory of induction eludes the problem simply because it is impossible to set up the difficulty, if one accepts the characterization of induction supplied by the material theory. An induction is licensed by its material postulate. But since the material postulate is itself another fact within our science, there is little mystery in what licenses it. It is licensed by other inductions based on other material postulates; and those in turn are licensed by further inductions based on further material postulates. This is a regress. But it is not the philosopher's fanciful fiction of an ascent through meta-inductions, meta-meta-inductions, and so on. Rather it is the real and mundane business of ordinary science that traces our justification for this fact in that fact; and that fact in some further fact; and so on. One might wonder how the resulting chains might end. As an empiricist I expect them to end in experiences. If there is a problem to be found one might expect it to be there. But a lot more needs to be said to establish a problem and what is said in setting up the traditional version of Hume's problem does not say it. For more discussion, see Norton (forthcoming, Section 6).

¹⁴ It is now apparent also why Bayesianism can have the broadest scope. It is a quite weak system that is able to license little until a large number of parameters are set. These are the prior probabilities and the likelihoods. The settings of these parameters can be used to encode a lot of factual information about particular domains. The material theory predicts that inductive inferences are particularized to the domains in which they arise; the settings of these parameters enables the Bayesian system to adapt itself to each domain with great flexibility.

References

- Achinstein, Peter (2001) *The Book of Evidence*. Oxford: Oxford University Press.
- Bacon, Francis (1620) *Novum Organum in Advancement of Learning: Novum Organum: New Atlantis*. Great Books of the Western World, Vol. 30. Chicago: University of Chicago Press, 1952.
- Cheeseman, Peter (1986) "Probability versus Fuzzy Reasoning," pp.85-102 in L. N. Kanal and J. F. Lemmer, eds., *Uncertainty in Artificial Intelligence*. Amsterdam: North Holland.
- De Finetti, Bruno (1937) "Foresight: Its Logical Laws, Its Subjective Sources". *Annales de l'Institut Henri Poincaré*, 7; reprinted *Studies in Subjective Probability*, eds. H. E. Kyburg and H. E. Smokler, New York: Wiley, pp. 95-157.
- Duhem, Pierre (1969) *To save the phenomena, an essay on the idea of physical theory from Plato to Galileo*. Trans. E. Doland and C. Maschler. Chicago, University of Chicago Press.
- Earman, John (1992) *Bayes or Bust: A Critical Examination of Bayesian Confirmation Theory*. Cambridge, MA: Bradford.
- Earman, John and Salmon, Wesley C. (1992) "The Confirmation of Scientific Theories," Ch.2 in M. H. Salmon et al. *Introduction to the Philosophy of Science*. rpt. Indianapolis: Hackett, 1999.
- Foster, Margueritte H. and Martin, Michael L. (1966) *Probability, Confirmation, and Simplicity: Readings in the Philosophy of Inductive Logic*. New York: The Odyssey Press.
- Giere, Ronald R. (1983) "Testing theoretical hypotheses," pp. 269-98 in J. Earman, ed., *Testing Scientific Theories: Minnesota Studies in the Philosophy of Science*, Vol. 10. Minneapolis: University of Minnesota Press.
- Gillies, Donald (2000) *Philosophical Theories of Probability*. London: Routledge.
- Glymour, Clark (1980), *Theory and Evidence*. Princeton: Princeton University Press.
- Goodman, Nelson (1983) *Fact, Fiction and Forecast* 4th ed., Cambridge, MA: Harvard University Press.
- Harman, Gilbert (1965) "Inference to the Best Explanation" *Philosophical Review*. 74 88-95.
- Hempel, Carl G. (1945) "Studies in the Logic of Confirmation" *Mind*, 54, pp. 1-26, 97-121; reprinted with changes, comments and Postscript (1964) in Carl G. Hempel, *Aspects of Scientific Explanation*. New York: Free Press, 1965, Ch.1.
- Howson, Colin and Urbach, Peter (1993) *Scientific Reasoning: The Bayesian Approach*. 2nd ed. La Salle, Illinois: Open Court.
- Janssen, Michel (2002) "COI Stories: Explanation and Evidence in the History of Science," *Perspectives on Science*, 10, pp. 457 – 522.
- Jardine, Nicholas (1984) *The birth of history and philosophy of science : Kepler's A defence of Tycho against Ursus, with essays on its provenance and significance*. Cambridge: Cambridge University Press.
- Joseph, H. W. B. (1916) *An Introduction to Logic*. 2nd ed. Oxford: Clarendon.
- Kelly, Kevin T. (1996) *The Logic of Reliable Inquiry*. New York: Oxford University Press.
- Lakatos, Imre (1970) "Falsification and the Methodology of Scientific Research Programmes," pp. 91- 196 in I. Lakatos and A. Musgrave, *Criticism and the Growth of Knowledge*. Cambridge: Cambridge University Press.

- Lipton, Peter (1991) *Inference to the Best Explanation*. London: Routledge.
- Mayo, Deborah (1996) *Error and the Growth of Experimental Knowledge*. Chicago: University of Chicago Press.
- McKeon, Richard (1941) *The Basic Works of Aristotle*. New York: Random House.
- Mill, John Stuart (1872) *A System of Logic: Ratiocinative and Inductive: Being a Connected View of the Principles of Evidence and the Methods of Scientific Investigation*. 8th ed. London: Longman, Green, and Co., 1916.
- Newton-Smith, William H. (2001) "Underdetermination of Theory by Data," pp. 532-36 in W. H. Newton-Smith, ed., *A Companion to the Philosophy of Science*. Blackwell.
- Norton, John D (1993), 'The Determination of Theory by Evidence: The Case for Quantum Discontinuity 1900-1915,' *Synthese*, 97 , 1-31.
- Norton, John D. (1994) "Theory of Random Propositions," *Erkenntnis*, 41, pp. 325-52.
- Norton, John D. (2000) "How We Know About Electrons," in Robert Nola and Howard Sankey (eds.) *After Popper, Kuhn and Feyerabend*, Kluwer, pp. 67-97.
- Norton, John D. (forthcoming) "A Material Theory of Induction," *Philosophy of Science*.
- Norton, John D. (manuscript a), "Must Evidence Underdetermine Theory?" First Notre Dame-Bielefeld Interdisciplinary Conference on Science and Values Zentrum für Interdisziplinäre Forschung, Universität Bielefeld, July 9-12, 2003.
philsci-archive.pitt.edu
- Norton, John D. (manuscript b), "The Formal Equivalence of Grue and Green and How It Undoes the New Riddle of Induction," philsci-archive.pitt.edu
- Popper, Karl, R. (1959) *Logic of Scientific Discovery*. London: Hutchinson.
- Salmon, Wesley (1967), *The Foundations of Scientific Inference*. Pittsburgh: University of Pittsburgh Press
- Salmon, Wesley C. (1984) *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.
- Savage, Leonard J. (1972) *The Foundations of Statistics*. New York: Dover.
- Shafer, Glenn (1976), *A Mathematical Theory of Evidence*. Princeton: Princeton Univ. Press.
- Stalker, Douglas (ed.) (1994) *Grue! The New Riddle of Induction*. Chicago and La Salle, IL: Open Court.