A Material Theory of Induction*

John D. Norton†‡

Contrary to formal theories of induction, I argue that there are no universal inductive inference schemas. The inductive inferences of science are grounded in matters of fact that hold only in particular domains, so that *all inductive inference is local*. Some are so localized as to defy familiar characterization. Since inductive inference schemas are underwritten by facts, we can assess and control the inductive risk taken in an induction by investigating the warrant for its underwriting facts. In learning more facts, we extend our inductive reach by supplying more localized inductive inference schemes. Since a material theory no longer separates the factual and schematic parts of an induction, it proves not to be vulnerable to Hume's problem of the justification of induction.

1. Introduction. There is a longstanding, unsolved problem associated with inductive inference as it is practiced in science. After two millennia of efforts, we have been unable to agree on the correct systematization of induction. There have always been many contenders. Some variant of Bayesianism now enjoys the leading position, although other schemes, such as inference to the best explanation, retain a considerable following. All this can change. In the late nineteenth century, some one hundred years after Bayes made his formula known, the leading systematization was the methods catalogued by Bacon, Herschel and, most precisely, Mill. This instability stands in strong contrast to deductive logic. The deductive syllogisms identified by Aristotle remain paradigms of deduction, with their very dreariness a mark of their unchallenged security. A comparable ancient contribution to inductive logic, induction by simple enumeration, has been a favored target of vilification for millennia. The problem is

*Received February 2003; revised May 2003.

†To contact the author write to Department of History and Philosophy of Science, University of Pittsburgh, Pittsburgh, PA 15260; e-mail: jdnorton@pitt.edu.

‡I thank Cristina Bicchieri, Phil Catton, John Earman, Kevin Kelly, Francis Longworth, Michela Massimi, Joke Meheus, Robert Nola, Wendy Parker, George Smith, and three anonymous referees for helpful discussion; and my special gratitude to Jim Bogen for significant help.

Philosophy of Science, 70 (October 2003) pp. 647–670. 0031-8248/2003/7004-0001$10.00

deepened by the extraordinary success of science at learning about our world through inductive inquiry. How is this success to be reconciled with our continued failure to agree on an explicit systematization of inductive inference?

It is high time for us to recognize that our failure to agree on a single systemization of inductive inference is not merely a temporary lacuna. It is here to stay. In this paper I will propose that we have failed, not because of lack of effort or imagination, but because we seek a goal that in principle cannot be found. My purpose is to develop an account of induction in which the failure becomes explicable and inevitable; and it will do this without denying the legitimacy of inductive inference. We have been misled, I believe, by the model of deductive logic into seeking an account of induction based on universal schemas. In its place I will develop an account of induction with no universal schemas. Instead inductive inferences will be seen as deriving their license from facts. These facts are the material of the inductions; hence it is a "material theory of induction." Particular facts in each domain license the inductive inferences admissible in that domain—hence the slogan: "All induction is local." My purpose is not to advocate any particular system of inductive inference. Indeed I will suggest that the competition between the well established systems is futile. Each can be used along with their attendant maxims on the best use of evidence, as long as we restrict their use to domains in which they are licensed by prevailing facts.

In Section 2, I will lay out the basic notions of a material theory of induction. Theories of induction must address an irresolvable tension between the universality and the successful functioning of some formal account of induction. The present literature favors universality over function. I urge that we can only secure successful functioning by forgoing universality and that this is achieved in a local, material theory of induction. In Section 3, I will review briefly the principal approaches to induction in the present literature in order to show that all depend ultimately on local matters of fact. In Section 4, I will illustrate how the imperfections of fit between existing schemas and actual inductions become greater as the relevant domain becomes narrower and suggest that some inductive inferences are best understood as individuals peculiar to a particular domain. In Section 5, I will review how a material theory directs that we control inductive risk by learning more facts since that both adds to the premises for our inductions and augments the inductive inference schemas locally applicable. In Section 6, I will argue that a material theory of induction eludes *The* Problem of Induction, in so far as the simple considerations that visit the problem on a formal theory fail to generate comparable difficulties for a material theory. Finally Section 7 contains concluding remarks.

2. The Material View.

2.1. The Problem of Enumerative Induction. Consider two formally identical inductive inferences:

Some samples of the element bismuth melt at 271°C. Therefore, all samples of the element bismuth melt at 271°C.	Some samples of wax melt at 91°C. Therefore, all samples of wax melt at 91°C.

The first is so secure that chemistry texts routinely report the melting points of elements on the basis of measurements on a few samples, relying on inductive inferences of exactly this type. The second is quite fragile, for reasons that are somewhat elusive, but certainly tied up with the fact that "wax" unlike "bismuth" is a generic name for a family of substances. Why is there such a difference? Mill ([1872] 1916, 205–206) found this question so troubling that he proclaimed "Whoever can answer this question knows more of the philosophy of logic than the wisest of the ancients and has solved the problem of induction."

2.2. Formal Theories of Induction. Let us look at attempts to answer the question in the presently dominant approach to induction embodied by what I call "formal theories." In them, the licit inductions are generated by supplying a universal template into whose slots factual content is inserted to generate licit inductions. The two inductions above are generated by suitable substitutions into the formal template of enumerative induction:

Some A's are B's.
Therefore, all A's are B's.

By formal theories, I intend something very broad. They are certainly not limited to accounts of induction within some formalized language or logic. The defining characteristic is just that the admissibility of an inductive inference is ultimately grounded in some universal template. Bayesian systems fall in this category. The probability calculus and Bayes' theorem provide a template into which degrees of belief are inserted.

Formal theories must respond to the problem of the melting points by insisting that the original schema of enumerative induction is elliptic and that extra conditions must be added to block misapplications. However it soon proves to be an insurmountable difficulty to find an augmentation that functions while still preserving the universality of the schema. For example, we might require that enumerative induction can only be carried out on A's that belong to a uniform totality. But without being able to mention particular facts about bismuth and wax, how are we to state a general condition that gives a viable, independent meaning to "uniform"?

We are reduced to the circularity of making them synonyms for "properties for which the enumerative induction schema works."

The problem of the melting points is closely related to Goodman's (1983) celebrated problem of "grue," where the same tension is revealed. A green, examined emerald confirms that all emeralds are green; and it also confirms that all emeralds are grue, where grue means green if examined prior to some future time T and blue otherwise. To explain why only the former confirmation is licit, Quine (1970) noted that green but not grue, is a natural kind term. However we cannot augment the schema for enumerative induction by requiring that A and B be natural kinds or natural kind terms without compromising its universality. For we have no universal account of natural kinds. In practice in science we determine which are the natural kinds by looking for the central terms of our current scientific theories. And even with that augmentation, we would still have no criterion for determining when an enumerative inductions is weak or strong; the strengths of the confirmation will still vary.

2.3. Material Theories of Induction. The natural solution to the problem of the melting points would seem to require explicit discussion of the differing properties of bismuth and wax. All samples of bismuth are uniform just in the property that determines their melting point, their elemental nature, but may well not be uniform in irrelevant properties such as their shapes or locations. Wax samples lack this uniformity in the relevant property, since "wax" is the generic name for various mixtures of hydrocarbons. A material theory of induction allows us to use such facts determine the differing strength of the inductions and rapidly resolves the problem.

In a material theory, the admissibility of an induction is ultimately traced back to a matter of fact, not to a universal schema. We are licensed to infer from the melting point of some samples of an element to the melting point of all samples by a fact about elements: their samples are generally uniform in their physical properties. So if we know the physical properties of one sample of the element, we have a license to infer that other samples will most likely have the same properties. The license does not come from the form of the inference, that we proceed from a "some . . ." to an "all. . . ." It comes from a fact relevant to the material of the induction. There are no corresponding facts for the induction on wax, so the formal similarity between the two inductions is a distraction.

In advocating a material theory of induction, my principal contention is that all induction is like this. *All inductions ultimately derive their licenses from facts pertinent to the matter of the induction*. I shall call these licensing facts the material postulate of the induction. They justify the induction, whether the inducing scientist is aware of them or not, just as the scientist may effect a valid deduction without explicitly knowing that it implements the disjunctive syllogism.

The material postulates determine the characters of the inductions in a material theory. They may certainly be truth conducive, as opposed to being merely pragmatically or instrumentally useful, as long as the material postulates are strong enough to support it. How each induction will be truth conducive will also depend on the material postulate and may well suffer a vagueness inherited from the present induction literature. Chemical elements are *generally* uniform in their physical properties, so the conclusion of the above induction is most likely true.[1] In this case, "most likely" denotes high frequency of truth among many cases of elements tested. Such frequentist readings will not always be possible.

2.4. All Induction is Local. There has been a long history of attempts to identify facts about the world that could underwrite induction. Best known is Mill's ([1872] 1916, book 3, chap. 3) "axiom of the uniformity of the course of nature." Russell ([1912] 1932, chap. 6) defined the principle of induction in terms of the probability of continuation of a repeated association between "a thing of a certain sort A" and "a thing of certain sort B." Russell's later (1948, part 6, chap. 9, 490–91) expansion has five postulates that include a quite specific "postulate of spatio-temporal continuity in causal lines" which "den[ies] 'action at a distance'.[2]"

All these efforts fall to the problem already seen, an irresolvable tension between universality and successful functioning. On the one hand, if they are general enough to be universal and still true, the axioms or principles become vague, vacuous, or circular. A principle of uniformity must limit the extent of the uniformity posited. For the world is simply not uniform in all but a few specially selected aspects and those uniformities are generally distinguished as laws of nature. So, unless it introduces these specific facts or laws, an effort to formulate the limitation can only gesture vaguely that such uniformities exist. Any attempt to characterize them further would require introducing specific facts or laws that would violate universality. On the other hand, if the axiom or principle is to serve its function of licensing induction, it must introduce these specific facts and forfeit universality. So Russell ends up denying action at a distance. If his account is to cover all induction, we must conclude that induction is impossible in any universe hosting action at a distance.

1. Why "generally"? Some elements, such as sulfur, have different allotropic forms with different melting points.

2. For another approach, see Keynes' (1921, chap. 23) presumption of "independent variety."

Because of these difficulties, the present material theory of induction is based on the supposition that the material postulates obtain only in specific domains; that is, facts that obtain "locally." As a result, inductive inference schemas will only ever be licensed locally.

The notion of a material theory has many precursors in the literature. Perhaps the most important is Goodman's "grue," which induced Hempel (see his 1964 Postscript in Hempel [1945] 1965, 51) to doubt the viability of a fully syntactic account of induction (see also Stalker 1994). My own path to the view is through the literature on demonstrative induction (see Section 5 below), in which inductions are reconfigured as deductive inferences with suppressed premises.

3. A Little Survey of Induction.

3.1. Universality Versus Successful Functioning. My principal argument for a local material theory of induction is that no inductive inference schema can be both universal and function successfully. As we saw in the case of enumerative induction above, to secure successful functioning we must forgo universality and adopt schemas that obtain only locally, under the license of locally obtaining facts. My purpose in this section is to complete the argument. I will try to show that this tension arises throughout the induction literature and make plausible that there are no successfully functioning schemas or systems of induction that do not in turn depend upon local matters of fact. While it will be impossible for me to address every system in the literature, I believe I can give reasonable coverage by recognizing that virtually all the systems can be fitted into one of three broad families, with all members of the same family exploiting the same inductive principle.

3.2. First Family: Inductive Generalization. All members of this family are based on the principle that an instance confirms the generalization. The best known and original form is enumerative induction, in which an instance of an A that is B confirms the generalization that all A's are B's. Hempel's (1945) satisfaction criterion of confirmation provides a logically robust definition of an instance in the context of first order predicate logic, while still exploiting the same idea that the instance confirms the generalization. A serious shortcoming of enumerative induction is that it merely licenses inferences from "some . . ." to "all . . ." and that is far too narrow for many applications in science. Augmented forms couple enumerative induction to schemes for introducing theoretical terms. The most important of these arise in Mill's ([1872] 1916, book 3, chap. 7) methods. For example, Mill's "joint method of agreement and disagreement" begins with our finding cases in which an antecedent always contains A if the consequent has B (necessity) and the antecedent always fails to have A if the consequent fails to have B (sufficiency). The joint method uses enumerative induction to infer that this necessity and sufficiency of A for B obtains not just in the cases investigated, but in all cases; and it licenses us to interpret this necessity and sufficiency as causation, so that A is the cause of B or an indispensable part of it. Glymour's (1980) bootstrap confirmation relation allows far greater liberty in the rules that can be used to introduce theoretical terms. The bootstrap confirmation relation requires that an instance of the hypothesis being supported is to be deduced from the evidence. In deducing it, we are free to use any part of the theory in question to introduce theoretical terms. In the context of Newtonian mechanics, we may use force = mass × acceleration to replace acceleration terms by force terms, so that evidence concerning accelerations can yield an instance of hypothesis pertaining to forces. In both these extensions, the basic inductive principle remains the same: a "some . . ." is replaced by an "all. . . ."

Since all these accounts still employ the same principle that the instance confirms the generalization, there is little to add to the discussion of Section 2 above. Whether an instance does confirm the generalization depends on the material postulate of the relevant domain; it will also determine the strength of the confirmation.

3.3. Second Family: Hypothetical Induction. Under the rubric of hypothetico-deductive confirmation, the basic principle of this family is that the ability of a theory, possibly with auxiliary assumptions, to entail the evidence is a mark of its truth. Its use is venerable, tracing back at least to the "saving of phenomena" in ancient astronomy. The basic principle by itself is actually rarely used alone because it is quite indiscriminate. According to it a trivial piece of evidence, that there is a triangle whose angles sum to 180°, is not just evidence for the Euclidean geometry that entails it but for the entire edifice of the science that may contain the geometry, including all of modern biology, geology, and astrophysics.

Expanding on such problems, we can quickly see that anything but an utterly trivial version of hypothetico-deductive confirmation requires additional restrictions that are factual. For example it allows that A&B is confirmed by A. The relation is non-trivial just in so far as it accords support to B, the second conjunct of A&B. But it should only do that if there is some relation of dependence between A and B. So A might be the report that we have measured the sum of the angles of a particular triangle and found it to be 180°; and B might be Freudian psychology. We would certainly not expect this A to lend support to this B, even though A&B does entail B. But it would support B if B were Euclid's fifth postulate of the parallels. For then A and B are related in that they are both parts of Euclidean geometry. More carefully, the relevance might be recovered from the assumption that the geometry is homogeneous, that is, of constant curvature. The presumption of a homogenous geometry functions as a material postulate in so far as it licenses inference from the value reported for the sum of the angles of a triangle, to the particular geometry prevailing, and then to the properties of parallel lines. Since we lack a general, formal account of the relation of dependence between A and B, a material fact decides whether the relation obtains and whether A confirms B.

There have been many attempts to tame the indiscriminateness of hypothetico-deductivism. They provide an interesting study in the tension between universality and successful functioning. In so far as they succeed in bringing proper functioning to the scheme, they do so by making the scheme dependent on particular facts, so that universality is compromised. These attempts can be grouped according to their broad strategy.

3.3.1. Exclusionary Accounts. There are infinitely many competing hypotheses able to entail the evidence. The most direct way to exclude all but a favored one is to demonstrate additionally that the falsity of the favored hypothesis entails the falsity of the evidence. That is rarely possible, since this additional result is very strong. It is equivalent (by contraposition) to the evidence deductively entailing the hypothesis. However our inductive worries would be essentially eliminated if we could show that the falsity of the hypothesis makes very likely the falsity of the evidence; or more briefly, that if the hypothesis weren't true we very likely wouldn't have gotten the evidence. Many accounts seek to augment the hypothetico-deductive scheme by demonstrating this additional result. As we shall see in the examples to follow, the demonstration of the additional results depends upon particular facts prevailing in the relevant domain. These facts function as the material postulate that licenses the induction in accord with the material theory.

The most straightforward of these accounts is modeled after traditional error statistical analysis; see Giere (1983) and, for a more thorough account, Mayo (1996). In a controlled study, any systematic difference in outcome between test and control group is attributed to the treatment applied to the test group. Typical study designs randomize over the two groups so that the only systematic difference in composition between the two groups is, most probably, the treatment which is then most likely responsible for differing outcomes. Generalizing these sorts of canonical examples, Mayo (1996, chap. 6) calls a test severe if the procedure is very unlikely to be passed if the hypothesis is false. Clearly passing a severe test is a strong license for an hypothesis. The license derives directly from the facts that make the test severe. These can be facts about the randomizing of test and control group, so that the "very likely" will be given through the physical probabilities of the randomizer. Or they may be vaguer. Replacing the bulb is a good test of the hypothesis that my lamp fails to operate because of a burnt out bulb, since it is very unlikely for another defect to be remedied by replacing the bulb. In this case, the probability judgments are vaguer, but still grounded in common experience of the world. That experience licenses the inference.

Facts that are far more difficult to identify underwrite another version of this augmentation. In arguments to a common cause (Salmon 1984, chap. 8) or in common origin inferences (Janssen manuscript) the hypothesis is inferred from the evidence on the strength of our conviction that it would be an astonishing coincidence if the evidence obtained without the hypothesis also being true. So early last century Perrin showed that there were roughly a dozen different methods of experimentally measuring Avogadro's number N, which gives the size of atoms.[3] That they all gave approximately the same value would be astonishing if there weren't atoms. Similarly, in developing his special theory of relativity, Einstein found it remarkable that our physical theories required material processes to contract, dilate, and more in just the perfect concert that made impossible any determination of the aether's presumed absolute state of rest. The astonishing coincidence could be eradicated if we conclude that all these processes are responding to the same background space and time that lacked an absolute state of rest.

What underwrites these inferences are elusive but widely shared judgments over what would happen if the relevant hypotheses were false. If there weren't atoms, we would expect some sort of continuum theory to prevail and we generally agree that these sorts of theories would not reproduce Perrin's experimental results. Similarly if the structure of space and time did harbor an absolute state of rest after all, we would expect that some physical process in space and time would reveal it. These are factual presumptions about the realm of possibility and they are the material postulates that underwrite the inferences.[4]

3.3.2. Simplicity. The most obvious and perennially popular augmentation of the hypothetico-deductive scheme uses the notion of simplicity

3. My contribution to this literature is to note this same inference form in Thomson's multiple measurements of the electron's mass to charge ratio. These experiments would not be expected to give the same value if his cathode rays did not in reality consist of discrete electrons (Norton 2000).

4. A closely related approach is Whewell's consilience of inductions. We are convinced of the correctness of Newton's physics since it does justice to both celestial and terrestrial realms, a coincidence that would be astonishing were Newton's physics false.

(see for example, Foster and Martin 1966, part 3) While many hypotheses or theories may deductively entail the evidence, in the augmented scheme we are licensed to accord inductive support only to the simplest of them. In my view, our decisions as to what is simple or simpler depend essentially upon the facts or laws that we believe to prevail. These facts dictate which theoretical structures may be used and the appeal to simplicity is really an attempt to avoid introducing theoretical structures unsuited to the physical reality governed by those facts or laws. That is, these facts or laws function as material postulates that license what we may or may not infer. Appeals to simplicity in the context of confirmation in science are really indirect appeals to facts presumed prevailing in the relevant domain.

This can be seen quite clearly in the most popular example of the use of simplicity in confirmation theory, curve fitting. When the evidence is presented as finitely many points on a sheet of graph paper and we find that a linear and a quadratic equation can be made to fit pretty much equally well, we routinely infer to the linear equation on the grounds that it is simpler. Our choice is licensed by facts tacitly assumed to obtain in the relevant domain.

To see this, first note that the applicability of the procedure depends upon getting certain assumptions right. An equation expressing a law is linear only with certain choices of variables. We can create an equation for the law that uses any function that strikes our fancy by the simple expedient of rescaling one of the original variables by the inverse of that function. Moreover there is nothing especially hallowed about the common choice of linear, quadratic, cubic, quartic, etc. as the natural sequence of functions of increasing complexity.

So picking the simplest curve can only make sense evidentially if we make the right choices for the variables and family of functions. And we make those choices correctly if we think that the variables and function hierarchy selected somehow map onto the basic physical reality at hand. So we plot observed planetary positions against the variable time and not the day of the week or middle initial of the observing astronomer. If our system is one describing growth, we would quickly look to fitting exponential functions since they are the functions that figure in the laws governing growth. In that case, we would certainly prefer an exponential curve over, say, a quintic curve that fitted as well, even though the exponential curve corresponds to a polynomial with infinitely many terms. Similarly if our system represents any sort of cyclic process, we would quickly look to sine and cosine functions, whose polynomial expansions have infinitely many terms, even though a fifth order polynomial might give us an equation with as many hills and valleys.

In short, we have no universal scheme or universal formal rules that define what is simpler or simplest. In so far as we are able, we choose the variables and functions appropriate to the facts that we believe prevail. These facts are the material postulates that license inference to the simplest curve.

3.3.3. Abduction: Inference to the Best Explanation. In this approach, we do not just require that the evidence be entailed by the hypothesis or theory it confirms; the evidence must also be explained by it (Harman 1965; Lipton 1991). The practice of abduction is straightforward. In each domain we are supplied with or may even create a repertoire of explanatory resources. We then choose among them for the one that explains the evidence best. So Hubble in the late 1920s observed that light from distant galaxies is shifted to the red, with the shift proportional to the distance. Galactic physics supplies several candidate explanations. The red shift may arise from a velocity of recession; or it may arise from the gravitational slowing of the distant galaxy's temporal processes, in accord with the general theory of relativity. The explanatory repertoire is limited. One might have expected it to include a process in which an interstellar medium reddens the light, just as a red filter might color a spotlight. Such a resource is not in the repertoire since chemistry supplies no medium that uniformly slides the spectral lines of the galactic light along the frequency scale. We routinely choose the velocity of recession as the explanation, since it explains the best, accommodating the linear dependence of red shift on distance to a uniform motion of expansion of the galaxies.

Examples in which the facts are uncertain underscore how the prevailing facts determine what may be inferred. Consider, for example, a controlled study of the healing efficacy of prayer. A theist would readily accept divine intervention as the best explanation of a positive result in the study. An atheist however would conjecture some as yet unnoticed flaw in the experimental design as the best explanation and may even deny that divine intervention would count as an explanation at all. The difference depends fully on their differences over the facts that prevail. An analogous example would be the differing appraisal of controlled studies on telepathy by parapsychologists and skeptics. Or in another example, astronomers and astrologers will differ markedly on how celestial bodies may figure as explanations of terrestrial phenomena. And modern day astronomers would no doubt not avail themselves of Newton's explanatory repertoire. He proposed ([1692] 1957, first letter) that the planets have just the right velocities to produce stable orbits since they were "impressed by an intelligent Agent."

The important point is that the facts prevailing in the relevant domain dictate what can count as an explanation for the evidence. Indeed that something explains at all is itself recovered by direct inspection of the case at hand and not by demonstrating conformity to some externally supplied template for what counts as a good explanation. We recognized that a velocity of recession is the best explanation of the galactic red shift without any thought to precisely what we might mean by saying that it explains the red shift. Did it explain by displaying causes, by subsumption under a covering law, by unifying, or by displaying statistically relevant factors?

So inferences to the best explanation are licensed by facts pertinent to the local domain that supply us explanatory resources. These facts express our expectations about the processes prevailing with the better explanations deemed more likely. The clearest case is when we take explanations to be displaying causes. The pertinent facts state what causes we expect to be active. We generally deem the better explanation to be the one that invokes the more likely case of fewer independent causes. So we are licensed to infer to the best explanation as the most likely.

3.3.4. Reliabilism. In a reliabilist account, we are licensed to believe hypotheses or theories produced by a reliable method. We routinely accept, for example, the diagnoses of expert car mechanics or physicians on the principle that the method they have used to arrive at their diagnoses is reliable, although we or even they may not understand the method. We might denounce an hypothesis as *ad hoc* merely because it was not produced by the right method. Reliabilists propose that science, properly practiced, uses reliable methods and that is why we can believe its products. I will consider such accounts as providing the augmentation needed to tame the indiscriminateness of hypothetico-deductivism, for that is a common use.[5]

We can see immediately that the material postulate that underwrites our inference in accepting these results is just our belief that the method is reliable. That belief is not only about the method. It is also about the world. For it incorporates the belief that the world is such that these methods can work reliably. In an uncooperative world, no method can succeed. We would properly hold little hope for a method touted as reliably finding a system to beat the house in a casino, for the casino sets up its games to make this impossible. Or we should have little hope in the prognostications of an entrail reader, no matter how expertly the best methods of entrail reading have been followed. The real world does not admit prediction through reading of entrails, except perhaps for the health of the flock from which the sacrificed animal was drawn.

5. In principle, however, hypotheses conjectured as the method proceeds need not be strong enough to entail the evidence and intermediate hypotheses need not even be consistent with it. For a general development of reliabilism that extends well beyond the framework of hypothetico-deductivism, in the context of formal learning theory, see Kelly 1996.

The method used may be quite explicit, in which case its reliability admits further analysis, either through other accounts of induction or directly through the facts that underwrite them. Or the reliability may be the practical end of the analysis, as is the case with expert diagnosticians. They do use explicit methods, codified in manuals. However these are supplemented by the experts' experience and in ways that may not admit explicit analysis.

The best known reliabilist account is Popper's (1959) falsificationism. According to it, scientists subject present theories and hypotheses to severe test by comparing their predictions with experience. If they are falsified, scientists are led to conjecture new theories or hypotheses and the latest products of this process are regarded as "well corroborated." While Popper has vigorously insisted (e.g. 33) that this notion of corroboration is not confirmation, I follow the tradition that holds that Popper's account simply fails to bear close enough resemblance to scientific practice if corroboration does not contain a license for belief, with better corroboration yielding a stronger license (see Salmon 1981). Popper does not give much account of the details of the method. The process of conceiving the new hypothesis is explicitly relegated to psychology and the inclination to take any philosophical interest in it disavowed as "psychologism" (31–32). So we have only our confidence in the scientist's creative powers to assure us that the new hypothesis does not introduce more problems than it solves. Lakatos (1970) gives a more elaborate accounting of a falsification driven method in his "methodology of scientific research programs." Decisions as to which research program a scientist ought to follow are dictated by guidelines deriving from such considerations as success at novel prediction. That we have any warrant to believe the products of such methods is licensed by factual matters: Is the method reliable? Are scientists actually following those methods properly? And more deeply: Is the world such that these methods can eventually succeed if followed?

The methods described are incomplete—they fall far short of an algorithm that can be followed mechanically. This once again reflects the tension of universality and successful functioning in the facts that underwrite induction. Perhaps we can eventually construct a full and complete account of a reliable method that works universally. Yet well before this extraordinarily optimistic goal is achieved, if we are so lucky as to have any complete methods at all, they will be narrowly specialized to quite particular domains, sacrificing universality for successful functioning.

3.4. Third Family: Probabilistic Accounts. In these accounts, degrees of belief are represented by magnitudes and the import of new evidence is to rearrange these magnitudes according to a definite calculus, usually the probability calculus. The inspiration for the family came from the discovery that the stochastic properties of physical systems could be represented by a calculus of probabilities. In Bayesian confirmation theory (Howson and Urbach 1989), this calculus is used to represent and update beliefs. Since it posits that our degrees of belief should conform to the same calculus as governs stochastic processes in the physical world, Bayesianism's least troublesome application arises when our beliefs pertain to stochastic processes. So, if a coin toss has a physical probability of 1/2 of yielding a head (henceforth a "chance" of 1/2), then our degree of belief in the outcome of a head ought to be that same 1/2. In this case, the material facts that license the ascribing of probabilistic degrees of belief are obvious: they are simply the physical properties of the coin toss system that generated the chances. They readily license truth conducive induction. The chance of at least one head in a run of ten independent tosses is more than 0.99. So if our degrees of belief conform to the chances, we will assign probability of least 0.99 to what will be the true outcome in over 99% of runs.

What makes matters more complicated is that Bayesianism asserts that the same notion of degrees of belief and the same calculus should be applied in all cases, including those in which no stochastic processes deliver convenient chances. So our prior beliefs in big bang and steady state cosmologies are probabilities measured on the same scale as our belief that a coin will come up heads or tails and they follow the same dynamics.

It is not too hard to conceive physical systems whose factual properties would preclude degrees of belief conforming to the probability calculus.[6] The steady state theory of cosmology posits that hydrogen atoms materialize randomly and uniformly in an infinite space. How are we to conform our beliefs to that fact for some hydrogen atom? We require a uniform probability distribution over all space. Traditionally that is achieved with an improper prior that assigns equal finite probability to equal finite volumes of space, at the cost of giving up the axiom of probability theory that requires the sum of all probabilities to be unity. In this case the sum is infinite (see Jeffreys 1961, §3.1).[7] In a slightly more complicated example,

6. We can also find systems for which our physical laws leave properties indeterminate and provide no probabilities for the various possible values. The best known example is the initial selection of matter density in the early universe within standard big bang cosmology. A Newtonian supertask system can also spontaneously excite with the theory giving no probabilities for the time or magnitude of the excitation. See Alper et al. 2000.

7. Or one might target another of the standard axioms of probability theory, countable additivity. That axiom allows us to sum the zero probability of appearance in each of the infinitely many individual cubic miles of space to recover a zero probability for appearance somewhere in all of space. Give up this axiom and we can retain a zero probability for appearance in any individual cubic mile of space, but a unit probability for appearance somewhere.

we might imagine the hydrogen atom constrained to materialize along some unit interval of space, [0,1]. The chances of the appearance in any subinterval would be given by the natural measure that, for example, assigns to the interval [1/4, 3/4] chance 1/2 = 3/4 − 1/4. Consider any subset of points in [0,1]. It is a reasonable question to ask what our belief is that the hydrogen atom appears at a point in the subset. However, if we require that our degrees of belief are probabilities conforming to the chances, then there will be infinitely many subsets for which no probability can be defined. These are the nonmeasurables of measure theory (Halmos 1950, §16).

There is another way in which we can see how material facts condition the content of Bayesianism and it does not call up esoterica like nonmeasurable sets. Bayesianism is vacuous until we ascribe some meaning to the probabilities central to it. Until then, they are just mathematical parameters. Each way of ascribing meaning brings factual presumptions with it. One popular way is to interpret the probabilities in terms of actions and recover their properties from our preferences. The best known of these is the Dutch book argument (De Finetti [1937] 1964). It is presumed that a degree of belief p in some outcome A makes us indifferent to which side of a bet we take in which winning $(1-p) is gained if A obtains and $p is lost if A fails. Given that we accept all bets of this form, it is quickly shown that we would accept a combination of bets (a "Dutch book") that force a sure loss, unless our degrees of belief p conform to the probability calculus. The relevant material facts are the rules for how we choose to convert degrees of belief into acceptable wagers, the existence of a framework for realizing the wagers and our preference for avoiding Dutch books. It is easy to see that these material facts would obtain in some contexts, such as if we are making wagers on horses at a racetrack. Even then, the details of the calculus that our degrees of belief would be constrained to follow are quite sensitive to the details of the rules. For example, if we decide that we are never indifferent to which side of a wager we accept, the avoidance of a Dutch book can no longer force additivity of our degrees of belief. There will be other contexts in which it is hard to see how these material facts could be made to obtain. How might our beliefs in big bang versus steady state cosmology be converted into the relevant sorts of wagers? Perhaps if one is a scientist choosing a research career, that choice might have the flavor of a wager. But one can still have beliefs on the theories when one has no personal stake at all in which is true and no interest in betting on it.

To the extent that the facts prevailing in a domain do not support realization of degrees of belief in bets, then avoidance of a Dutch book fails to constrain those degrees of belief. The obtaining of these facts, however, does *not* comprise a material postulate strong enough to support truth conducive induction (unlike the earlier case of material facts pertaining to stochastic processes). Rather they merely license inductions that are pragmatically useful in so far as they generate beliefs that will not support actions that may bring sure losses.

Nevertheless Bayesian confirmation theory seems to be the most general of all the systems currently in favor. The reason is that it is a rather weak system. We get very little from it until we specify what might need to be a quite large number of conditional probabilities (the likelihoods) and these are determined by the factual properties of the relevant system. Each such specification yields a mini-inductive logic adapted to the facts of the relevant domain. So we might ask if the hypothesis H that all swans are white is confirmed by the evidence E that this swan is white in the natural sense that $P(H|E) > P(H)$. Even though we already know the likelihood $P(E|H) = 1$, we cannot answer *at all* until we specify the likelihood $P(E|\text{-}H)$, which requires us to judge how likely it is to find a white swan if not all swans are white. These likelihoods, which determine the mini-inductive logic, are in turn fixed by our prior probability distribution, since, for example, $P(E|\text{-}H) = P(E\&\text{-}H)/P(\text{-}H)$. So we have the curious result that a mythical prior probability distribution, formed in advance of the incorporation of any evidence, decides how any evidence we may subsequently encounter will alter our beliefs.

4. Inductions Too Local to Categorize. In the little survey above, we can see that the more universal the scope of an inductive inference schema, the less its strength. Unaugmented enumerative induction and hypothetico-deductivism can assure us only of weak support at best. As we narrow the domain of application of an inductive inference schema, it can be grounded in more specific matters of fact and can supply stronger support. At the same time an imperfection of fit will arise between it and the small repertoire of inductive inference schemas recognized in the literature. Thus we should expect cases of inductive inferences that prove to be too hard to characterize, while at the same time we are quite sure of their strength from merely inspecting the particular case.

The clearest examples of such inductions arise when a major advance in the science brings a major enhancement to our inductive powers. In such cases, the inductive inference schemas supplied by the methodology literature remain essentially unchanged. The enhancement in our inductive reach can be most naturally attributed to the new material postulates made available in the relevant domain.

Prior to Lavoisier and the establishment of modern chemistry in the late eighteenth and early nineteenth century, it was quite hard to know what sorts of properties of substances were likely to be the same across all samples. Learning which substances were elements and compounds and which mixtures dramatically improved our inferential abilities. To know that something was an element or compound brought a license to infer that all its samples were most likely alike in physical properties. We also learned when to infer that our failure to bring about a transformation might merely be due to our failure to find the right methods or when it could be generalized to an impossibility. From a chemical perspective there were no barriers to transforming inorganic matter into organic matter; it was a case of the former. But our failure to transform lead into gold was a case of the latter. While these all fit the form of enumerative inductions (with some failing) they have become so completely modified by the particular chemical facts in the domain that the characterization has almost no practical value.

More examples: Once we learned in the 1920s that the nature and properties of the elements are due to the quantum properties of electrons trapped by atomic nuclei, we had a much stronger basis for knowing which unfilled spaces in the periodic table might really coincide with undiscovered chemical elements, where the table might be expanded and where no such expansion would be possible; and we secured a greater ability to decide when a new substance with certain stable properties might be a new element. After Newton showed us the gravitational forces that act between celestial bodies, we were given a new prescription for inferring to causes. All we needed to show was that some effect could be generated within the repertoire of Newton's system and we could infer to its reality. So Newton himself showed that the moon's gravitational attraction caused our tides and that comets were deflected in their motion about the sun by the force of gravity from the sun. The scheme even licensed inferences to new bodies. The planet Neptune was discovered in the nineteenth century by working back to the location of an undiscovered body that could cause perturbations in the planet Uranus' motion.

In these cases, the added inferential power that comes from knowing more does not come from delivery of some new schema. In the cases above, it is even hard to know what to call the schemas. The inference to the moon's gravity as cause of the tides or to a new planet is not just simply finding an hypothesis that saves the phenomena. It has to do it in the right way. One might be tempted to talk of best explanations, common causes or consiliences. None quite capture the strength of the inference; some inferences to best explanations or common causes can be weak and some strong. The clearest explication of what that right way amounts to is just a local fact: the hypotheses do it in accord with the repertoire of Newtonian gravitation theory. Our confidence in Newton's theory underwrites the strength of the induction.

5. The Control of Inductive Risk.

5.1. Strategies in a Formal and Material Theory. In inductive inference we take an inductive risk: the danger that we may accept or accord high belief to a result that turns out to be false. In science we seek to reduce this inductive risk as much as possible. The strategies for controlling inductive risk are different according to a formal theory of induction or a material theory.

According to a formal theory, we approach the problem in two ways. First we seek to amass as much evidence as possible. The better the evidence the stronger will be our inductive inferences. Second we seek to expand the inductive inference schemas available to us. While this second approach might seem promising, with the notable exception of continuing work in statistics, there has been relatively little work by scientists devoted to expanding our repertoire of inductive methods.

According to the material theory of this paper, this lacuna is not so surprising. The two approaches to controlling inductive risk cannot be separated. We reduce our exposure to inductive risk by collecting more evidence. At the very same time, exactly because we learn more from the new evidence, we also augment our inductive schemas. For according to the material theory, all these schemas obtain only locally and are ultimately anchored in the facts of the domain. Crudely, the more we know, the better we can infer inductively. The result is that scientists do not need to pay so much attention explicitly to inductive inference. As we saw in the examples of Section 4, with each major advance in science has come a major advance in our inductive powers. The mere fact of learning more will augment their inductive powers automatically.

5.2. The Portability and Localization of Inductive Risk. The above examples also illustrate a common dynamic in our efforts to control inductive risk. We start with an induction that uses some fairly general schema—enumerative induction or hypothetical induction. The inferences are risky because the generic forms of the schemas are known to be unreliable. We localize the induction to a particular domain, whose material postulates licenses the induction far more securely. In the examples, with some risk we generalized the physical properties of one sample of a substance to all. When we recognize that substance is an element, we now have recourse to the known constancy of elemental properties to underwrite the inference securely. Because of the correlation of tides with the position of the moon, we hypothesize it as the cause of the tides. Localizing the process to gravitation theory and drawing on the resources of Newton's theory, we become much more confident that the hypothesis is correct.

This dynamic is a transporting of inductive risk from a schema to a fact, the relevant material postulate. This portability affords an important means of assessing and controlling inductive risk. As long as the inductive risk resides within the schema, we must assess it through a highly problematic judgment of the overall reliability of the relevant schema. We have little chance of coming to a clear judgment let alone determining how to reduce the risk. However once the risk is relocated in a material postulate in some local domain, our assessment of the inductive risk will depend in large measure on our confidence in the material postulate. If the inductive risk is great, we now also have a program for reducing it. We should seek more evidence relevant to the material postulate and perhaps even modify the material postulate in the light of the evidence. The result will be a more secure induction.

In short, we can control inductive risk by converting schematic risk into presumptive risk, since the latter can be more accurately assessed and reduced.

5.3. Demonstrative Induction as a Limiting Case. If we can reduce inductive risk by transporting it from the schemas into the material postulates, might we eliminate it entirely? We can, in a limiting case in which the material postulate and the evidence taken together deductively entail the hypothesis at issue. There is no trickery in this limiting case. In effect we are just discovering that we have already taken the relevant inductive risk elsewhere in our investigations when we accepted the material postulate. As a result we do not need to take it again.

This form of inference has entered the literature under many names, such as demonstrative induction, eliminative induction or Newtonian deduction from the phenomena. The latter term is appropriate since demonstrative inductions arise quite naturally in Newton's analysis of planetary motions (See Smith 2002; Harper 2002). For example, Newton knew that the supposition that planets are attracted by an inverse square law to the sun is sufficient to generate the stationary ellipses of the planetary orbits. But that the hypothesis saves the phenomena is not decisive. Might there be other force laws that also yield stationary ellipses? Newton showed that this was not possible. In Propositions 44 and 45 of his *Principia,* he showed that alternatives would fail to yield the stationary orbits observed. For the near circular orbits of the solar system, he showed any force law in $1/r^n$ (with r the distance to the sun) would yield a rotation of the axis of the apsides, the points of farthest and nearest approach to the sun, from which the value of n could be read. The evidence of the fixity of this axis conjoined with the relevant propositions deductively entails that $n = 2$ and that the force law is an inverse square law.

There are numerous other examples in the literature. For example, following Planck's 1900 work, it was recognized that the hypothesis of quantization of energy would save the phenomena of the distribution of energy over different frequencies in heat radiation. But is that enough to force us to accept this hypothesis so fundamentally at odds with classical physics? Shortly after, in the early 1910s, Ehrenfest and Poincaré showed that we had to accept it. The extra presumptions needed to make the inference deductive were ones already accepted: essentially that thermal systems are really just systems with many degrees of freedom acting in the most probable way. From the relevant phenomena, they now showed one could deduce the quantization of energy (see Norton 1993; and for an example in Bohr's work, Norton 2000).

6. The Problem of Induction Eluded?

6.1. The Analog of Hume's Problem for the Material Theory. In order to learn a fact by induction, the material theory says that we must already know a fact, the material postulate that licenses the induction. Does some vicious circularity or harmful regress lurk here? One will quickly recognize this concern as the analog of a familiar problem for formal theories of induction, re-expressed in the framework of a material theory of induction. It is just *The* Problem of Induction, that most celebrated of philosophical problems traditionally attributed to Hume. We shall see that Hume's problem can be set up quite easily for a formal theory, since a formal theory separates factual content from formal schemes. I will argue that the absence of this separation in a material theory results in the same considerations failing to generate a comparable problem for a material theory.

In the usual context, the problem of induction asserts (Salmon 1967, 11) that there can be no justification of induction. A deductive justification would violate its inductive character; an inductive justification would either be circular or trigger an infinite regress. To generate the analogous problem for a material theory, we consider the material postulates that justify inductions in the two cases. Analogous to the deductive justification of induction is the use of a material postulate that is a universal truth known *a priori*. That justification fails since such a postulate would violate the locality of induction. Analogous to the inductive justification of induction is a material postulate that is a contingent fact. If that fact is the same fact as licensed by the induction, then we have an obvious circularity. If it is a different fact, then we trigger a regress. But, I shall urge, the regress is neither infinite nor demonstrably harmful. The analogy is summarized in the table:

TABLE 1. HUME'S PROBLEM OF INDUCTION FOR A FORMAL THEORY AND ITS ANALOG IN A LOCAL MATERIAL THEORY

Formal Theory		Material Theory	
Justify an inductive inference schema . . .	Diagnosis	Analogously, justify an induction to a fact . . .	Diagnosis
. . . by a deductive argument?	Fails. Violation of inductive character of induction.	. . . by a material postulate that is a universal truth known *a priori*?	Fails. Violates local character of induction.
. . . by displaying many successful instances of the schema and applying the *same* inductive schema?	Fails. Circular.	. . . by using the *same* fact as the material postulate?	Fails. Circular.
. . . by displaying many successful instances of the schema and applying a *different* inductive schema?	Fails. Infinite regress of fanciful meta- and meta-meta-inductions is triggered.	. . . by using a *different* fact as the material postulate?	A regress is triggered through a sequence of justifying facts, but it is neither infinite nor demonstrably harmful.

6.2. Comparison. The problem is immediate and serious for a formal theory. Consider the justification offered in the last line of the table. We justify an inductive inference schema by displaying many successful instances of it and performing a meta-induction on them using a different inductive inference schema. This first step is already a fanciful proposal, since we do not actually carry out such meta-inductions on inductions scattered through our science. But it is just the beginning of an infinite regress of inductions. If we can reconcile ourselves to the first meta-induction, then we must face a second of even greater ambition: a meta-meta-induction on the meta-inductions. This must continue indefinitely. Each meta-induction will be of broader scope and more remote from any inductions anyone does or could do. The sequence has no termination. We face a fatal infinite regress.

When we transport the argumentation used to set up the problem of induction to the material theory, it no longer forces the same sort of difficulty.[8] While the first two justifications of Table 1 are obviously problematic, the third is not. In it, we induce a fact with an induction that is

8. I thank Jim Bogen for making me see this.

grounded by the facts of a material postulate; these latter facts are justified by inductions that are in turn grounded in the facts of other material postulates; and those facts are justified by inductions grounded in other facts; and so on. The regress described here is far from the fanciful meta-meta-meta-inductions remote from actual inductive practice required by a formal theory. It merely describes the routine inductive explorations in science. Facts are inductively grounded in other facts; and those in yet other facts; and so on. As we trace back the justifications of justifications of inductions, we are simply engaged in the repeated exercise of displaying the reasons for why we believe this or that fact within our sciences.

What remains an open question is exactly how the resulting chains (or, more likely, branching trees) will terminate and whether the terminations are troublesome. As long as that remains unclear, these considerations have failed to establish a serious problem in the material theory analogous to Hume's problem. And it does remain unclear. It is possible that serious problems could arise in termination. In principle the chains could end in some sort of circularity, although such circularity was not displayed in any of the examples above. It is also possible that the chains have benign termination. They may just terminate in brute facts of experience that do not need further justification, so that an infinite regress is avoided. Or, more modestly, they may terminate in brute facts of experience augmented by prosaic facts whose acceptance lies outside the concerns of philosophy of science—for example, that our experiences are not fabricated by a malicious, deceiving demon. Perhaps we might doubt that a single such brute fact is rich enough to license a substantial induction. But we should not expect that of a single brute fact. It is more reasonable to expect that enough of them, careful woven together through many smaller inductions, would eventually license something grander. A decision for or against must await the ever elusive clarification of the notion of brute facts of experience and of whether the notion even makes sense.[9]

The closest we have come to a fatal difficulty in the material theory is a regress whose end is undecided but with the real possibility of benign termination: a fatal difficulty has not been forced. Contrast this with the analogous outcome in a formal theory: a regress whose *beginning* is problematic and whose end, an assured infinity, is disastrous.

In sum, the simple considerations that visit *The* Problem of Induction on formal theories fail to generate a comparable difficulty for a material

9. I reject the simple argument that such brute facts are always singular and that no collection of singular facts can license a universal. The problem is that canonical singular facts—"the ball is red"—already presuppose universal knowledge. In this case it resides in the recognition that the thing is a ball and that its surface is red, thereby admitting recognition of commonality with potentially infinite classes of objects and colors.

theory. What makes the difference is that the material theory does not separate facts from inductive inference schema. This separation diverted a formal theory into tracing the justification of an induction back through an infinite hierarchy of meta- and meta-meta-inductions that no-one actually does or could complete. In analogous circumstances, the material theory merely traces the justification of an induction back through chains of licensing facts whose justification is part of the regular practice of science.

7. Conclusion. My purpose in this paper has not been to advocate any particular scheme of inductive inference from the many that compete in the literature of philosophy of science. Rather I want to suggest that they are all admissible in the right context and to try to explain why we have such a proliferation of them in enduring conflict. I have urged that we resolve the intractable tension between the universality and the successful functioning of an inductive inference schema by forgoing universality and adopting a material theory of induction. In such a theory, the facts that prevail in each local domain in science license inductive inference schemas that are peculiar to that domain. We justify the inductive inferences of the domain by reference to these facts and not by passing through them to universal inductive inference schemas. I have tried to show how the existing schemas for inductive inference all require some local facts for their justification. I have also suggested that any schema with pretensions of universality will fit actual inductions imperfectly and that the fit will become worse as we proceed to narrower domains and the facts licensing the inductions become more specialized. This, I believe, explains a curious phenomenon in science. We can be quite sure of a result in science as long as we look at the particulars of the result and the evidence that supports it, but we often end up struggling to explain by means of standard inductive inference schemas how the evidence can yield that strength of support. Finally I have suggested that the material theory gives new options for assessing and controlling inductive risk: we investigate the warrant for the material postulate; and the theory tells us that merely learning more facts can extend our inductive reach by supplying more local inductive inference schemas.

REFERENCES

Alper, Joseph S., Mark Bridger, John Earman, and John D. Norton (2000), "What is a Newtonian System? The Failure of Energy Conservation and Determinism in Supertasks", *Synthese,* 124: 281–293.

De Finetti, Bruno ([1937] 1964), "Foresight: Its Logical Laws, Its Subjective Sources", in Henry E. Kyburg and Howard E. Smokler (eds.) *Studies in Subjective Probability.* New York:Wiley, 95–157. Originally published in *Annales de l'Institut Henri Poincaré,* 7.

Foster, Margueritte H., and Michael L. Martin (1966), *Probability, Confirmation, and Simplicity: Readings in the Philosophy of Inductive Logic*. New York: The Odyssey Press.

Giere, Ronald R. (1983), "Testing Theoretical Hypotheses", in John Earman (ed.), *Testing Scientific Theories,* Minnesota Studies in the Philosophy of Science, vol. 10. Minneapolis: University of Minnesota Press, 269–298.

Glymour, Clark (1980), *Theory and Evidence*. Princeton, NJ: Princeton University Press.

Goodman, Nelson (1983), *Fact, Fiction and Forecast,* 4th ed. Cambridge, MA: Harvard University Press.

Halmos, Paul R. (1950), *Measure Theory*. Princeton, NJ: Van Nostrand.

Harman, Gilbert (1965), "Inference to the Best Explanation", *Philosophical Review*. 74: 88–95.

Harper, William (2002), "Newton's argument for universal gravitation" in I. Bernard Cohen and George E. Smith (eds.), *The Cambridge Companion to Newton*. Cambridge: Cambridge University Press, chap. 5.

Hempel, Carl G. ([1945] 1965), "Studies in the Logic of Confirmation", reprinted with changes, comments, and Postscript (1964) in Carl G. Hempel, *Aspects of Scientific Explanation*. New York: Free Press, chap. 1, 3–51. Originally published in *Mind* 54: 1–26, 97–121.

Howson, Colin, and Peter Urbach (1989), *Scientific Reasoning: The Bayesian Approach*. La Salle, IL: Open Court.

Janssen, Michel (2002), "COI Stories: Explanation and Evidence in the History of Science", *Perspectives on Science* 10: 457–522.

Jeffreys, Harold (1961), *Theory of Probability,* 3rd ed. Oxford: Clarendon Press.

Kelly, Kevin T. (1996), *The Logic of Reliable Inquiry*. New York: Oxford University Press.

Keynes, John Maynard (1921), *A Treatise on Probability*. London: MacMillan.

Lakatos, Imre (1970), "Falsification and the Methodology of Scientific Research Programmes", in Imre Lakatos and Alan Musgrave (eds.), *Criticism and the Growth of Knowledge*. Cambridge: Cambridge University Press, 91–196.

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.

Mill, John Stuart ([1872] 1916), *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.

Newton, Isaac ([1692] 1957), "Four Letters to Richard Bentley" in Milton K. Munitz (ed.), *Theories of the Universe*. New York: Free Press, 211–219.

Norton, John D. (1993), "The Determination of Theory by Evidence: The Case for Quantum Discontinuity 1900–1915", *Synthese* 97: 1–31.

——— (2000), "How We Know about Electrons", in Robert Nola and Howard Sankey (eds.), *After Popper, Kuhn and Feyerabend*. Boston: Kluwer, 67–97.

Popper, Karl R. (1959), *Logic of Scientific Discovery*. London: Hutchinson.

Quine, W. V. O. (1970), "Natural Kinds" in Nicholas Rescher (ed.), *Essays in Honor of Carl Hempel*. Dordrecht: Reidel, 1–23. Reprinted in Stalker (1994), 41–56.

Russell, Bertrand ([1912] 1932), *The Problems of Philosophy*. Reprint, London: Thornton Butterworth.

——— (1948), *Human Knowledge: Its Scope and Limits*. New York: Simon and Schuster.

Salmon, Wesley (1967), *The Foundations of Scientific Inference*. Pittsburgh: University of Pittsburgh Press.

——— (1981), "Rational Prediction", in Adolf Grünbaum and Wesley C. Salmon (eds.), *The Limitations of Deductivism*. Berkeley: University of California Press, chap. 5.

——— (1984), *Scientific Explanation and the Causal Structure of the World*. Princeton, NJ: Princeton University Press.

Smith, George E. (2002), "The Methodology of the Principia", in I. Bernard Cohen and George E. Smith (eds.), *The Cambridge Companion to Newton*. Cambridge: Cambridge University Press, chap. 4.

Stalker, Douglas (1994), *Grue! The New Riddle of Induction*. La Salle, IL: Open Court.