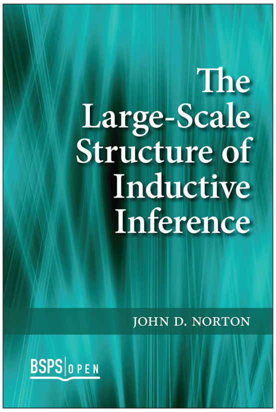

THE LARGE-SCALE STRUCTURE OF INDUCTIVE INFERENCE

John D. Norton

ISBN 978-1-77385-541-7

Acknowledgement: *We acknowledge the wording around open access used by Australian publisher,* **re.press**, *and thank them for giving us permission to adapt their wording to our policy* http://www.re-press.org

Introduction

1. The Project of This Volume

According to the material theory of induction, inductive inferences or relations of inductive support are not warranted in a way familiar from accounts of deductive logic. They are not warranted by conformity with some universally applicable schema or template. Rather, each is warranted by background facts peculiar to the domain in which the inference arises. This idea was developed in my earlier monograph, *The Material Theory of Induction* (Norton 2021). A key provision of the theory is that the warranting facts must be facts: that is, truths of the domain. If we seek to sustain an inductive inference by appealing to some warranting proposition in the domain that is false, then we commit the inductive analog of a fallacy. The error is comparable to the deductive fallacy of affirming the consequent as if it were a valid deductive schema.

That warrants must be factual truths places a special burden on us when we assess the inductive inferences or relations of inductive support among the propositions of some science. To establish support fully, we must also establish the truth of the warranting propositions used. Since these warranting propositions are also contingencies of the domain, establishing their truth requires further inductive inferences. Thus, any claim that some particular item of evidence inductively supports some other proposition in a theory is not self-contained. To be sustained to the fullest extent, the truth of these further warranting facts must also be established. Since those facts themselves are contingent propositions, we must establish their truth with still further inductive inferences or relations of inductive support, and we must show that those inductive inferences in turn are warranted by further facts.

And so on. All claims of inductive support, in effect, are claims that concern a large network of contingent propositions within the science of interest and, commonly, extending beyond it.

These considerations define the project of this work. Individual claims of inductive support must be made within a larger ecology of relations of inductive support. How is this larger ecology configured? What is the large-scale structure of inductive inference? What are its problems? Can a cogent account be supplied for it? My goal in this work is to answer these questions.

Some might find this entanglement of inductive support with a larger inductive ecology disquieting and might want to retreat to formal approaches to escape it. Formal approaches that use universal schemas might appear to have an advantage. They can assess the cogency of an inductive inference without engaging a larger ecology. An inductive argument from analogy just has to show that it conforms to the relevant schema. A claim of probabilistic support might just have to show that the associated probabilities relate by Bayes' theorem.

This advantage is illusory. According to the material theory of induction, it is dangerous to assume that each formal schema can be applied unconditionally everywhere. It exposes users to a significant risk of inductive fallacies if the schemas are applied in domains that lack a material warrant. The common remedy by formalists is tacitly to limit the application of the schemas to where they are thought somehow to be appropriate. The remedy is poor since decisions on applicability depend on hunches and intuitions. Here material theorists have the advantage. The question of which inference forms are applicable where is decided by an explicit analysis of the prevailing facts.

Again one might think that a better way to treat the large-scale structure of inductive support is mathematical. We merely need to identify the calculus that applies at this large scale. Questions about the large-scale structure would be answered mathematically by theorems in the calculus. Bayesians in the philosophy of science might believe that the probability calculus already does just this.

Hopes for some universal calculus of inductive inference fail and provably so. In recent work (Norton 2019, 2021, Chapter 12), I have shown the incompleteness of all calculi of inductive inference that meet some minimal conditions. Any such calculus will fail to discern nontrivially the inductive import of any body of evidence unless the computation is supplemented by inductive content supplied externally. The familiar example is that Bayesian analysis always requires prior probabilities. Their stipulation is antecedent to the application of Bayes' theorem, yet their content exerts a strong influence on the outcome of the computations. Efforts have failed to supply Bayes' theorem with vacuous priors that exert no such influence. This incompleteness is not limited to the probability calculus. A form of it will arise in any calculus meeting minimal conditions.

In its briefest form, the answer supplied by the material theory of induction to the question of the large-scale structure of inductive support is this: relations of inductive support within a mature science form a massively entangled network without any clear hierarchical structure. Quine (1951, 39–40), in his celebrated "Two Dogmas of Empiricism," presented a similar structure for beliefs. However, his structure was variously a "fabric" and a "field of force" and later a "web of belief." Its key attribute is its elasticity. A conflict with experience, according to this picture, can always be accommodated. The internal connections, he supposed, are so elastic that there are many ways to do this. This supposition has been responsible for much philosophical mischief. It has encouraged the idea that evidence, even in great measure, is unable to determine the propositions of a science. This indeterminacy is incompatible with our routine experiences of mature science and is not established by Quine's analysis. The elasticity results from reliance on a naive and inadequately weak hypothetico-deductive approach to inductive inference.[1]

The account developed in this volume differs sharply from Quine's (1951) supposition of elasticity. The relations of inductive support in a mature science are better imagined as strong steel cables, not elastic threads. They are connected and interconnected in such a variety of ways that the integrity of the entire structure is threatened if an anomalous experience arises. The affirmation that some ordinary machines can be combined to produce a perpetual motion machine would overturn mechanics. Or consider the discovery of a new mineral not constituted by atoms or not compounded of elements found in the periodic table. It would destabilize chemistry and, after that, the quantum theory that underpins the atomic character of matter and the uniqueness of the elements in the periodic table. Evolutionary theory would fail to accommodate a new species of living beings that spontaneously appears fully formed without any past history of development. The structure

1 Or so I argue (Norton 2008).

of inductive support of mature sciences is not elastic but rigid. A break in one place propagates with revolutionary import far into the structure.

In this volume, I explore and examine this structure. The first chapter is a brief development of the material theory of induction. It does not replace the lengthier elaboration of the theory in *The Material Theory of Induction* (Norton 2021). However, for readers interested in the issues raised in the present work, it will serve well enough as a point of first contact.

Subsequent chapters are divided into two parts. The first part presents general propositions in the philosophy of science concerning the large-scale structure of inductive inference or inductive support. The second part presents historical case studies that provide detailed illustrations of the main claims of the first part and are the source of many of its claims.

2. Part I: General Claims and Arguments

In Chapter 2, I advance four claims, whose support and elaboration occupy the remainder of the text.

1. Relations of inductive support have a nonhierarchical structure.
2. Hypotheses, initially without known support, are used to erect nonhierarchical structures.
3. Locally deductive relations of support can be combined to produce an inductive totality.
4. There are self-supporting inductive structures.

The first claim renounces the idea that inductive support is hierarchical, structured by generality. In this renounced picture, propositions in a science are supported inductively just by propositions of lesser generality. We then would be able to trace a pathway of inductive support from the lowest levels closer to experience, gradually ascending the hierarchy of generality to the most general propositions of the science. The actuality is that relations of inductive support in real science fail to respect any such hierarchy. They cross over in many complicated ways. The very idea of a hierarchy of generality is sustainable only in a crude way, if at all.

The second claim pertains to the practices needed to identify these tangled inductive structures. In the early stages of the development of a new science, inductive inferences commonly can proceed only if we make use of warranting assumptions for which we do not yet have inductive support. They are introduced as hypotheses, and their use is provisional. Their use comes with an obligation to secure their proper inductive support in subsequent investigations. Should that obligation not be met, the original claims of inductive support fail. This role attributed to hypotheses is *not* the traditional role given to them in accounts of hypothetico-deductive confirmation. In this latter case, the hypotheses themselves are confirmed by their success at entailing evidence. Here the hypotheses mediate in establishing inductive support for other propositions. The hypotheses themselves must accrue support by other means in another stage of investigation.

The third claim asserts that it is possible to combine deductive relations of support to produce an overall relation of support that is inductive in character. This is a possibility that, in the abstract, seems to be impossible. Yet, as the examples show, it arises commonly in the actual practice of science. If it can be achieved, then it is a construction to be prized for its reduction in inductive risk. The more familiar construction involves intersecting relations of inductive support that are combined to produce an overall inductive import. An inductive risk is taken, first, in accepting each component relation or inductive inference and, second, in accepting their combined import. When the component inductive relations of support are replaced by deductive relation, that first inductive risk is eliminated.

Finally, the fourth claim is a thesis of completeness. That many inductive inferences are warranted materially is undeniable, or at least so I think after working through the many examples in *The Material Theory of Induction* (Norton 2021). If one is eager to retain general schemas, then it is tempting to suppose that these examples display only a part of the full inductive story. Materially warranted inductive inferences or relations of support alone, one might want to assert, are not enough to sustain all of a science inductively. A full accounting must include general schemas or general rules in some form. This fourth claim asserts otherwise. It is possible for materially warranted propositions to form structures such that every proposition in the structure is inductively supported, without the need for general schemas or other devices outside the material theory of induction.

This completeness is already a corollary of the arguments given for the material theory in Chapter 2 of *The Material Theory of Induction* (Norton 2021) and repeated more briefly in Chapter 1 below. Any general schema must factually expand in some way on the premises supplied to it. This expansion can only be sustained in domains hospitable to the means of the expansion. Any such expansion can fail if the facts of the domain are such as to oppose it. The fact of that hospitality, in the most general terms, is the warranting fact of the inductive inference or relation of inductive support. This argument defeats every attempt to assert the existence of some universal inductive rule. There can be none that escapes it.

A simpler picture does not use universal inductive rules. Each individual proposition of a mature science is inductively well supported. If we are willing to undertake the task of tracing it, we can display the form that support takes and its material character. This is true of each of the propositions of a mature science, taken individually. Their totality is the full material account of the inductive support of the mature science. Nothing further is needed, for no proposition has been left without an account of its inductive support.

These four claims in turn raise further issues that need to be addressed. Relations of inductive support cross over one another in a myriad of ways. Tracing along the pathways of support, we routinely find circles that bring us back to our starting point. Philosophers, brought up in fear of vicious circularity, mistakenly find the mere existence of such circles automatically disqualifying for any system. In Chapter 3, I argue that this disqualification is hasty and mistaken. There are circles throughout our sciences. We routinely consider populations in which the rate of growth of the population is proportional to the size of the population. This is a benign circle of self-reference. It is merely the most convenient definition of exponential growth. When a circularity is uncovered, there can be no default supposition of a systemic failure. Instead, we have a positive obligation to demonstrate that a circularity is harmful, if we seek to represent it as such. Is the circularity vicious and thus leads to a contradiction? Or does it lead to an underdetermination of theories? I argue that the circularities in inductive relations of support within mature scientific theories do neither. They are benign.

In Chapter 4, I address a related issue. Mature sciences, it has been asserted, are inductively self-supporting. The evidence for them is sufficient to sustain relations of inductive support such that every proposition in the science is supported. That leaves open a troubling possibility. Might it be that there are multiple such sciences for a given body of evidence? Then the bearing of evidence would not be univocal, no matter how rich and varied the evidence. Might this be the harm that circularities bring? I argue otherwise in the chapter. Mature sciences are uniquely supported by their evidence. There is only one periodic table of elements supported by the evidence in chemistry and so on for the central claims of mature sciences.

This uniqueness arises from the empirical character of science. Any alternative is only a real alternative if it differs in some factual assertion. Since all such assertions are open to empirical testing, competition among alternatives is transient, if only the evidence that can decide among them is pursued. The material character of inductive inference adds a mechanism that destabilizes any competition. If one theory in the competition gains a small advantage, then the facts thereby secured can serve as warrants for further inductive inferences supporting the theory. The effect is that the advantage of the ascending theory is amplified. When the investigations continue, this amplification is repeated, at the repeated costs of its competitors. If the process continues long enough, then it ends with one theory prevailing over all of its competitors. It is this instability that promotes the uniqueness of mature sciences.

Circularities are a distinctive feature of coherentist accounts of justification. We might hope, as I did originally, that there would be results already developed there of use to the material theory. The differences between the two systems are so great that, it turns out, these expectations are not met. In Chapter 5, I explore these differences. The coherentist account is offered as an alternative to fundamentalist accounts of justification. Coherentists must eschew the fundamentalist supposition that some beliefs are justified primitively by the world. The material theory has no such obligation. It takes observations and experiences of the world to be the foundations upon which inductive structures are built. For coherentists, beliefs are justified by their inclusion in a coherent system. The judgment is essentially global. There is something similar in the material theory. Strong inductive support for a proposition ultimately does depend on the larger-scale integrity of the relations of inductive support. However, that integrity arises from the composition of many individual relations of support. Each of the propositions in the structure must be well supported inductively, and considerable effort is expended in establishing each such individual relation of support. Finally, coherentist justifications concern relations among beliefs: that is, within cognitive states. The material theory is concerned with mind- and belief-independent relations of inductive support among propositions that assert some factual condition in the world.

In Chapter 6, I describe how the material theory of induction dissolves the classic problem of induction. I provide a short history of the problem and show that the problem of induction is specifically a problem for accounts of induction based on universal schemas. Its dissolution by the material theory involves no exotic legerdemain. The material theory of induction does not posit universal schemas. It follows that the problem of induction cannot be set up in it. It is dissolved. Although this claim of dissolution has already attracted considerable attention, it has come with the mistaken claim that the material was devised specifically to solve the problem of induction. As I have related on several occasions, that is not the history of it.[2] My concern is that the claims of the material theory — at both the local scale and the large scale — should be evaluated as an attempt to understand inductive inference better. That can be done independently of whether the theory dissolves the problem of induction. If it does not dissolve that problem, then the failure merely puts it in good company with all of the other failed attempts. The material theory's other results still stand.

Although the material theory's dissolution of the problem of induction is straightforward, a common reaction is to treat it like other claimed solutions to the problem. Under scrutiny, these other solutions prove to depend on unfounded, hidden assumptions, comparable in import to those that produced the problem originally. This reflexive reaction leads to the supposition that the problem must reappear in the material theory in some way in the mutual dependencies of inductive support. The unmet challenge of this reflexive reaction has been to find a way that the problem of induction reappears. Perhaps circularities in the structure are harmful; perhaps there is a fatal regress to warranting propositions of ever greater generality; perhaps, if our starting point is meager, then we have no warranting hypotheses that would allow inductive inferences to be initiated. All of these suppositions fail to identify a problem for the material theory. There is little need to argue the point in great detail since securing the theory against such objections is undertaken in earlier chapters. The theory's circularities are benign, as I argue in Chapter 3. A fatal regress to warranting propositions of ever greater generality requires the presumption of a hierarchical structure that, as I argue in Chapter 2, is

2 My first paper on the material theory (Norton 2003) was already in a complete first draft when Jim Bogen pointed out the possibility of a dissolution of the problem of induction. I added an imperfect sketch of that dissolution as a later section of the paper.

not present in the material theory. Finally, there is no difficulty starting the inductive project. When warranting premises are missing, they are introduced provisionally as hypotheses.

3. Part II: Historical Case Studies

The second part of this volume presents a set of case studies within the history of science. They are detailed and reflect my commitment that an analysis of inductive inferences should be responsible to what actually happens in science. Here the analysis differs from much of what is found in the philosophical literature on inductive inference. There the analysis suffers from adaptation to a few oversimplified examples. We might infer from the observation that some crows are black the conclusion that all crows are black. But such inferences, analyzed in isolation, are oversimplified caricatures of the much more sophisticated inductive inferences of real science. An account designed just to accommodate such oversimplified examples is destined to be woefully oversimplified itself.

Formal accounts of inductive inference in the philosophical literature face the same problem. An erudite formal analysis, no matter how technically clever, is only as good as the assumptions on which it is based. The inductive practice of real science is complicated and messy. Formal systems, if they are to be amenable to mathematical analysis, must be based on a few simple axioms. When they are naive or oversimplified, so, inevitably, is the analysis. These failures are easily overlooked since, commonly, formal accounts are developed without close attention to the actual inductive practices in science. When a formally pretty system is proposed, it is easy to be distracted by the ingenuity of the technical details and beguiled by the lure of the abstract formal puzzles that they pose.

This work takes seriously the obligation to connect its general claims with the actualities of the sciences. It does this by melding general claims in the philosophy of science with detailed historical studies of science. This practice embodies a conception of what it is to do the history and philosophy of science. Theses in the philosophy of science must withstand scrutiny in the history of science. That much is widely accepted as an abstract principle. It is much less widely practiced. The reverse relation is more interesting. I have found repeatedly that investigations in the history of science are a fertile means of identifying powerful and interesting theses in the philosophy of science. The scientists often face daunting inductive challenges. Their ingenuity in meeting the challenges far outstrips the imaginings of the philosophers of science, concerned only with ruminations on abstract principles and ideas. Careful attention to the history can yield ideas that otherwise would not emerge from mere armchair reflection.

Chapters 7 to 14 present case studies selected, I must admit, simply because they are episodes that interest me and, I suspect, will prove to be fertile in supplying general theses for the first part. In almost all of the cases, we find relations of inductive support crossing over one another in a way that violates a hierarchy of generality. That is one of the most important facts provided by the studies. The individual studies typically add extra points of special interest.

In Chapter 7, I recount Hubble's arguments in 1929 for his celebrated "Hubble's law." That law asserts that galaxies recede with a speed proportional to their distance from us. If one does not look at the details of his reporting, then it is all too easy to represent his analysis as a simple act of generalization. Hubble checked that the linear relation held for a sample of galaxies and then just generalized. A little attention to his paper of 1929 shows that his analysis was neither so simple nor that easy. Hubble had distance measurements only for roughly half of the galaxies in his data set. He needed maneuvers of great ingenuity to extend his law to all of the galaxies in his data set. They involved reasoning that inverted the order of inference that one would expect. In one part, they even employed the Hubble law itself as a premise.

In Chapter 8, I recount some of Newton's arguments for his inverse square law of gravity. Newton, we find, was adept at recovering inductive support for his claims by combining deductive relations. Such combinations figure in central portions of the evidential case that he made for his theory of universal gravitation. They arise in his Moon test, which argues for the identification of terrestrial gravity and the force that binds the Moon to the Earth; they arise again in the details of his analysis of the inverse square law of gravity and its relation to the elliptical orbits of the planets.

In Chapter 9, on atomic spectra, I show how the numerical rules governing the series of lines in the hydrogen emission spectrum are supported by multiple relations of inductive support that cross over one another in many ways. Under the warranting authority of Ritz's combination principle, the presence of some lines provided support for the presence of other lines, and entire infinite series of lines provided support for other entire infinite series of lines. A second crossing over of support occurred at a higher level. Ritz's combination principle provided general support for the newly emerging quantum theory. It was the observable manifestation of the fundamental electronic process of Bohr's quantum theory of the atom: the stepwise descent of an excited electron through the allowed orbits of the theory. Soon this relation of support was inverted. The more fully developed quantum theory both entailed the Ritz combination principle and could specify the empirically found circumstances in which it failed.

In Chapter 10, I provide another illustration of the crossing over of relations of support. It arises among two sets of propositions that date historical artifacts. In one set, datings are provided by traditional historical and archaeological methods. In the other, datings are provided by radiocarbon methods. There are uncertainties in both. Historical methods can err when they rely on meager or equivocal clues. Carbon dating can err if the historically varying levels of atmospheric carbon 14 are not accurately known. Then the baseline from which the carbon 14 decay started is uncertain. Each set can be used to correct and calibrate the other set. The calibration curve for historical levels of atmospheric carbon 14 was derived from historical dating methods, including, famously, the counting of tree rings in ancient bristle cone pines. Once well calibrated, carbon 14 dating can correct historical and archaeological datings of artifacts. When the two sets of propositions are in agreement, each mutually supports the other set.

In Chapter 11, I look at the history of the determination of the relative atomic weights of the elements. The task proved to be recalcitrant and strained the resources of chemists for roughly the first half of the nineteenth century. The difficulty was that, after Dalton's introduction of chemical atomism in 1808, chemists were trapped by an incompleteness in his atomic theory. The evidence that 8g of oxygen react with 1g of hydrogen to produce water does not tell us how many atoms of hydrogen combine with how many atoms of oxygen to form water. Was the ratio 1-1, 2-1, 1-2, and so on? We are left uncertain about whether the molecular formula for water is HO, H_2O, HO_2, or something else again. To eliminate the uncertainty, we also need to know the relative weight of each atom of hydrogen and oxygen. But we cannot know those relative weights until we know the correct molecular formulae for water and other related substances.

Chemists struggled for roughly half a century to overcome this incompleteness. Matters were settled only with Cannizzaro's results in 1858 and brought to the notice of chemists through a conference in 1860. Cannizzaro's results depended on a careful selection of fertile hypotheses to break the evidential circle in which Dalton was trapped. The best known is Avogadro's hypothesis on the numbers of molecules in equal volumes of gases. Applying this and other hypotheses to a wide array of elements and compounds, a unique set of atomic weights could be recovered. They emerged from a huge tangle of intersecting relations of support. There were so many that I can sample only a few in the chapter. They extend from intersecting relations of support among the molecular formulae of individual substances to mutual relations of support at the highest levels of abstract theory. The chemists found support for Avogadro's hypothesis in the new physics of the kinetic theory of gases. Conversely, the physicists found support for their new physics in the chemists' adherence to Avogadro's hypothesis.

In Chapter 12, I provide another illustration of the importance of hypotheses in enabling inductive investigations to proceed. Since antiquity, astronomers have sought to determine the distances to the Sun, Moon, and planets. Simple methods of geometric triangulation — called "parallax" when used astronomically — provided only meager results. The angles to be measured were too small for naked eye astronomy to resolve reliably. That changed when telescopic observations became possible in the seventeenth century. The task remained formidable. Attempts to use parallax for this purpose still called for major scientific expeditions as late as the eighteenth and nineteenth centuries.

These observations and simple geometry alone were not enough. Hypotheses were required to warrant inferences from the observations to the distances sought. Distances thus inferred remained provisional until independent support was provided for the hypotheses. Early hypotheses used in these investigations failed to meet the requirement. Ptolemy derived his estimates of the distances to the Sun, Moon, and planets using the hypothesis that space is filled with the spheres of his geocentric cosmology, packed together as closely as possible. His distance estimates collapsed when his geocentric cosmology failed to find the independent support needed. Reliable distance measurements were subsequently recovered only with the mediation of the Copernican hypothesis, in turn further supported by Newton's theory of universal gravitation. These hypotheses did accrue the requisite independent evidence.

The last two chapters provide examples of theories in competition. They are intended to illustrate the claims of the instability of inductive competitions described in Chapter 4. In Chapter 13, I examine the practice of dowsing. Miners in the Harz mountains of Germany in the sixteenth century believed that minerals underground can be detected by the deflections of a hazel twig. Over the centuries, dowsing migrated to the detection of underground water.

The competition recounted was between dowsers and their skeptical critics and how it turned to favor the skeptics. The case for dowsing was mostly secured anecdotally. It lay in repeated accounts of dowsing successes and even the mere existence of a profitable profession of dowsers. The critics were able eventually to challenge successfully the reliability of these accounts. The nineteenth-century identification of ideo-motor effects explained how dowsers erroneously might have come to believe that the effect was real. On the theoretical side, by the rudimentary standards set by the early theories of electric and magnetic attraction, it was plausible that underground minerals might exert an influence above ground. Over the centuries, the growth of theories of electricity and magnetism left no theoretical space for the mechanism of dowsing. The critics' successes in these two strands of phenomena and theory were mutually supporting and came at the cost of proponents of dowsing. By the early twentieth century, dowsing had been reduced to the status of a pseudoscience.

In Chapter 14, I recount a present-day case of systems of prediction in ongoing competition. I recount four systems, all of which are currently applied to predict the future movement of prices on the stock market. They are fundamental analysis, technical analysis ("chartists"), random walk/efficient market analysis, and fractal/scale-free analysis. The competition among the systems is lively. Proponents of each are aware of the competing systems and try to impugn them. I provide a sample of their disagreements. The guiding principles of each system are hypotheses in the sense of Chapter 2. They are proposed provisionally to enable prediction to proceed. However, none has been secured evidentially such that it has found universal acceptance. That follows from the persistence of the disagreements among the proponents of the individual systems. However, these hypotheses are mutually exclusive: at most, one can be true. The evidence that would single it out is available in abundance in the past history of trading on the stock market. Were this evidence to be pursued and evaluated without prejudice, the disputes would be resolved, and at most one system would prevail. Instead, however, we have the curious spectacle of proponents who refuse this task. The disagreement continues in full display, and we can continue to watch how each approach seeks to gain an evidential advantage over the others.

REFERENCES

Norton, John D. 2003. "A Material Theory of Induction." *Philosophy of Science* 70: 647–70.

———. 2008. "Must Evidence Underdetermine Theory?" In *The Challenge of the Social and the Pressure of Practice: Science and Values Revisited*, edited by Martin Carrier, Don Howard, and Janet Kourany, 17–44. Pittsburgh: University of Pittsburgh Press.

———. 2019. "A Demonstration of the Incompleteness of Calculi of Inductive Inference." *British Journal for the Philosophy of Science* 70: 1119–44.

———. 2021. *The Material Theory of Induction*. BSPS Open Series. Calgary: University of Calgary Press.

Quine, Willard Van Orman. 1951. "Two Dogmas of Empiricism." *Philosophical Review* 60: 20–43.