DRAFT

Chapter from a book, *The Material Theory of Induction,* now in preparation.

Prolog

John D. Norton
Department of History and Philosophy of Science
University of Pittsburgh
http://www.pitt.edu/~jdnorton

1. The Wonder of Science

Our best science tells us wonderful things. The cold and dark skies of our universe were not so long ago in their entirety in a state of unimaginably high energy and temperature. The detritus that exploded from it congealed into stars, planets and galaxies. These systems of celestial masses are in turn held together by a curvature of the geometry of space and time itself. On a most minute scale, the matter of these systems and the light they radiate consist of neither waves nor particles but a curious amalgam that is, at once, both and neither. The organisms that walk on one of these planets, complete with their intricate eyes and thinking brains, emerged incrementally from crude matter, in tiny steps over eons of time. They were shaped only by the fact that a small, random change in one organism might give it a slight advantage over its rivals. The design specification of these accumulated advantages is recorded and transmitted through the generations of the organisms by its encoding in hundreds of millions of base pairs of a chemical found in every cell of each organism.

These, and many more ideas of science like them, are extraordinary. Their contemplation must eventually overwhelm with wonder even the most curious and flexible of minds. Only the dullest of wit or the most soured of skeptics could resist their charms.

For me, there is a still greatest wonder. These ideas are not the inventions of writers of myth and fiction. They could not be so, for their content far outstrips our meager human imaginations. Rather they are the result of careful, painstaking, systematic investigations of nature, guided solely by inventive insight and cautious reasoning. They are discoveries. When these efforts go past the early speculative stages and succeed, their products are distinguished by a special relation with what we experience of the world. Those experiences provide the inductive support for successful science. They tell us that this is how the world is.

The explosive expansion of the universe is supported by the reddening of light from distant galaxies. That the curvature of the geometry of space and time keeps the planets in their orbits is supported by the most delicate measurements of slight anomalies in planetary motions. The curious quantum nature of matter in the small is supported by how light from excited gases is concentrated into just a few quite specific frequencies. The evolution of we humans from simpler organisms is supported by fossilized bones, whose chronology is recorded by their positions in layers of rock strata. The double spiral geometry of the molecules of deoxyribonucleic acid is supported by the patterns formed when X-rays diffract off material extracted from the nuclei of cells.

In all this, the essential relation is inductive support. It obtains between the propositions of science and those that express the evidence on which it rests. It enables us to assign an authority to the ideas of science that no other narrative can match. Without it, science becomes just another "way of knowing," to use a popular oxymoron of the skeptics. Without this relation, we do not know anything of the world. We "know" but do not know. Without it, the ideas of science are no better than the fanciful creation stories of primitive mythologies.

2. Where the Philosophy of Science Literature Falls Short

If we are to understand how science succeeds where these other narratives fail, we must understand how this relation of inductive support works. That is a core task for philosophy of science. Its efforts reside in the expansive literature on induction or inductive inference. The project of this book results from an enduring dissatisfaction with this literature.

There is no shortage of approaches in this literature. However, what is distinctive about these approaches is that they are fractured. There are many of them. They rise and fall with the generations and even with the particular philosopher consulted. Each has its successes and each has its failures. None, it seems to me, is by itself fully adequate to the task.

Loosely speaking, there two traditions.[1] One is qualitative and a few examples illustrate its pervasive problems. Evidence supports those hypotheses that, in various senses, generalize the evidence; or deductively entail the evidence; or explain the evidence; or provide a severe test of the evidence. Each case is troubled. There are so many ways one item of evidence can be generalized that most generalizations cannot be supported. Most applications of the simple scheme must fail. Similarly there are very many hypotheses that entail one item of evidence. The same problem arises. Most applications of this scheme will fail. The problem of proliferation is ameliorated if the hypothesis must not just entail the evidence but explain it. The meagerness of the gain is revealed when we realize that we have no general account of explanation precise enough to support a theory of inductive inference. The account rests ultimately on dubious intuitive judgments of what explains what and how well it does it. Severe testing requires a judgment that the evidence would likely not come about were the favored hypothesis false. To apply the scheme we must know what is likely in the case of this falsity. Excepting contrived situations like controlled studies, such judgments are at best speculative and at worst self-serving inventions.

The second tradition is quantitative. We assign a numerical measure to the support. The measure used almost universally is probability. The approach is, initially, appealing since we replace a vague "weakly supports" or "strongly supports" by precise numbers that must be combined by quite specific rules. Now we can calculate! My enthusiasm for this approach dampened when I found that its central theoretical tool, Bayes' theorem, has a voracious appetite for prior probabilities and likelihoods. The trouble is that their values must be specified by considerations outside the calculation itself. Prudent or malicious choices for their values, more than the niceties of mathematical theorems, control the final result. Worse, as this Bayesian approach ascended to the momentary dominance it presently enjoys in the literature, its analyses became more and more separated from real applications to inductive inference in the sciences. They have drifted towards self-contained exercise in recreational probability theory. That separation is disguised by tendentious labeling of terms. A calculation best adapted to the accumulated results of many coin tosses is represented as giving some sort of understanding of how the accumulation of intricate and diverse evidence in science can support a univocal result.

[1] This is a hasty dissection of an enormous literature. See Norton (2005) for a more careful dissection and categorization.

The situation has not been improved by a rash decision to conceive of the prior probabilities of Bayes' theorem subjectively, that is, as freely chosen opinions that can vary from person to person. For once one has let arbitrary opinion into the system, the probabilities cease to measure strengths of inductive support, but only some indissoluble amalgam of them with arbitrary opinion. These problems are not resolved but compounded with dubious analogies. We are told a fable of a punter at a racetrack making monetary bets with bookies who are determined to take every advantage possible. This epistemic situation is supposed sufficiently close to that of scientists weighing evidence for big bang cosmology or a neural basis for cognition that all should conform to the same principles of rationality.

3. The Material Approach

The upshot of these accumulated woes is that philosophy of science as a discipline cannot now offer those outside it a univocal account of inductive support. My goal in this book and in the larger program of research it embodies is to solve this problem. The clue to its solution is found in the observation that each of the accounts sketched above do work somewhere. If we are investigating controlled trials, then ideas about severe testing are apt. If we are interested in matching DNA from blood samples with that of accused offenders, then we can use Bayesian methods. When Einstein found that his new general theory of relativity "explained" (as he put it) the anomalous motion of Mercury, he could claim a wonderful "confirmation" (as he wrote) of his theory.

The clue in all this is that the application of the various approaches work when we add factual conditions that limit the domain in which they are to be applied. The stronger the factual restriction, the more successful the application. The material approach simply asks us to "take the limit." That is, what warrants the successful application of the particular inference is found *entirely* in the background factual conditions that delimit the domain of application.

This last assertion is the key idea of the material theory. It distinguishes it from all other approaches. They use the standard literature in deductive inference as the model for analyzing inductive inference. It provides them a formal model. According to it, we distinguish the good from the bad inferences by checking whether the candidate inference fits in its form with some universal template or schema. For example, take the inference

All men are mortal.

Therefore, some men are mortal.

This is a valid, deductive inference since it is derived from the universally applicable schema that I will call "*all-some"*:

All *A*'s are *B*.

Therefore, some *A*'s are *B*.

We are allowed to make any substitution for *A* and *B* and we are assured that what results will be a good inference in its form. The schema is universally applicable. Its use is not restricted, for example, to inferences about human mortality.

Since antiquity, philosophers have sought to recover similar schemas for inductive inference. The successes have always been partial. One of the earliest attempts was "enumerative induction":

Some *A*'s are *B*.

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

The trouble is all too clear. It will almost never work. With obvious substitutions, we might be happy to infer:

Some men are mortal.

Therefore, all men are mortal.

But we would be unhappy with almost every other variant of it, such as:

Some men are Greeks.

Therefore, all men are Greeks.

All of the approaches sketched briefly above lie within this formal tradition. If we just focus on simple examples like these, it becomes quite apparent that they must fail to have universal scope.

The schema *all-some* does have universal scope since it is fully self-contained. Its cogency derives completely from the meanings of the words "all" and "some." If someone doubts the cogency of the inferences it authorizes, we would gently inquire of them whether they understood the meaning of the words.

In contrast, enumerative induction is not self-contained. It can work, but only when we restrict the substitutions for *A* and *B* to terms hospitable to the induction. When *A* is "men," successful substitutions for *B* include biological properties like "is mortal," "is borne of a mother," "has a blood circulation system," and so on. That is, if we restrict the domain in which the scheme is applied, it can warrant good inferences. However its success is entirely dependent on the restriction. The facts comprising that restriction are the ultimate source of its warrant. They are biological facts about people. The inference is warranted, in the last analysis, because that is the way people are biologically.

Further, the inference is a good inference only in so far as the warranting facts are true. If science advances so far that we can create people entirely in the test tube from synthetic DNA without the need for a gestating mother, these facts would cease to be true and one of the inferences would become an inductive fallacy.

It is easy to see how these conclusions about inductive inference generalize. All inductive inferences lead to conclusions that go beyond what is necessitated logically by their premises. It follows that they are only good in so far as the inferences are carried out in domains that are factually hospitable to the inferences. The facts that make the domain hospitable are the facts that warrant the inference. Here it is helpful to remember that a commonplace of deductive inference is that that propositions can both state factual matters and also serve as warrants for deductive inference. The proposition "If *A* then *B*." is both a factual proposition and also a warrant that authorizes a deductive inference from *A* to *B*. The material theory asserts that, ultimately, this dual role for factual propositions is the only way that inductive inferences are warranted.

This applies even to Bayesian analysis, in so far as it has any ambitions of providing an account of inductive inference. It is true that the manipulations of Bayes' theorem itself are deductive inferences lying within the probability calculus. We deduce a value near unity for the probability of Newton's universal law of gravitation, conditioned on the motion of the sun's planets and their moons. An essential background fact is that these deductions are implemented in a domain in which distributions of inductive support are properly represented by probabilities. In the second half of this book, we shall explore domains in which this presumption fails.

These last considerations constitute the core of the material approach to inductive inference. It provides a single, unified approach that incorporates all the different approaches presently in the literature; or at least it incorporates them all in so far as they are sufficiently precisely defined to be viable in some domain.

Its core ideas can be encapsulated in some slogans: "All induction is local." This slogan reminds us that any regularity we may find among inductive inferences is restricted to some domain and dependent for its warrant on the particular facts that obtain there. Another slogan is

“There are no universal rules for inductive inference.” It reflects the core posit that the warrant of an inductive inference is not traced back, ultimately, to some universal schema, but to facts that obtain only locally.

If one hears only this slogan in isolation, one might mistake it for a skeptical thesis akin to Feyerabend’s notorious “anything goes.” That is very far from its import. It is merely a part of the relocating of the warrant of inductive inferences from rules to facts. The material theory does not seek to undermine inductive inference. It seeks to save it. For the formal approaches that dominate the literature have simply failed in their most important functions. None gives us a successful system, applicable universally, for discerning which are the good inductive inferences. None gives an account of why the inferences it does authorize are appropriate. This last failure stands in stark contrast with standard examples of deductive inference. Inferences warranted by the deductive schema *all-some* are good inferences simply in virtue of the meaning of “all” and “some.” These last considerations pose two problems that the material theory solves.

First, inference schemas in the present literature cannot be used universally. While their writings are curiously silent on the question, Bayesians will concede to me in conversation that their system does not apply everywhere. That invites the key questions of where are the limits and how we identify them. The material theory answers. One must locate the facts that can warrant the schema, Bayesian or otherwise. The schemas can be applied only in domains in which those facts obtain.

Second, merely stating an inference schema does not automatically make it a good one. In familiar deductive cases, we discern that they are good because of the meaning of the connectives. We cannot do the same for inductive schemas. Instead, the material theory tells us that certain inference schemas are good since they depend on factual matters in the domain of application. Biological predicates, like “is mortal” and “has a blood circulation system” are facts common to all people and that fact of commonality authorizes the inferences sketched earlier.

Adopting the material approach to inductive inference leads one to approach problems in inductive inference differently. There is no default scheme that can be applied mechanically and automatically. If one wants to employ some mode of inductive inference in some context, one must be able to supply positive reasons for why that mode is applicable in that circumstance. This applies especially to probabilistic inference. One should not assume by default that it always applies. If it is to be used in some domain, we have a positive obligation to provide the foundations for its applicability. Otherwise it cannot be used.

While this book is largely not concerned with beliefs (credences) as opposed to objective relations of inductive support, the moral carries over. There should not be a default presumption that credences are probabilities. If credences are to be represented as probabilities in some circumstance, then positive reasons must be given for why they are appropriate in that circumstance.

4. The Chapters

The chapters of this book are divided into two parts. The earlier Chapters 1-9 are devoted to laying out the basic ideas of the material theory and applying it to what are identified above as the qualitative approaches to inductive inference. The later Chapters 10-16 concerns quantitative approaches, most notably the probabilistic approaches of Bayesianism.

Chapter 1 states the basic propositions of the material theory of induction. The vehicle to develop them is Marie Curie's inference from the crystallographic properties of her sample of Radium Chloride to those of all possible samples. It is an instance of enumerative induction of breathtaking scope. It depends on the evidence of just a few specks of the only sample of Radium Chloride then known. This chapter also shows how the material theory can warrant successful inferences of this form, even if of breathtaking scope, by displaying the underlying facts that warrant them. In this case the pertinent fact is Haüy's principle. It lies at the core of extensive investigations of into the properties of crystals in the nineteenth century and solves the vexing problem of discerning just which of the many properties of crystals are projectable, that is, suitable for enumerative inductions.

Chapter 2 elaborates the argument stated briefly above in Section 3 that justifies the material theory of induction. The essential ideas of the justification are these. No extant formal scheme of inductive inference has proven to be applicable universally. The successes of all these schemes can be explained by the material facts within the restricted domains in which they succeed. Most importantly, inductive inference is by its nature ampliative. That means that its conclusions are logically stronger than its premises. Hence an inductive inference can only succeed in domains in which further background facts are hospitable to it. This chapter also poses the inductive puzzle “1, 3, 5, 7. What’s next?” The puzzle is, of course, insoluble nontrivially without some indication of the background facts that can serve to warrant an inductive inference that answer “what’s next?” The chapter reports the underappreciated and ingenious way Galileo solved the problem.

Subsequent Chapters 3 to 9 address specific rules and schemes proposed in the literature for inductive inference. The goal of these chapters is to show that, when these rules or schemes work, they do so because of identifiable background facts; and that they can only work in domains with such hospitable facts. We also find in each case that the apparent unity of application of the candidate rule survives only as long as we do not look too closely at the details of the examples. As we consider those details more thoroughly, we find the specific background facts taking on the primary burden of warranting the inferences. The original rule survives only as a superficial similarity among the examples.

In writing these chapters, I have tried as much as possible to use examples of inductive inference from real science. This literature can suffer when commonplace, non-scientific examples are used to guide our inductive inferences in science. The material theory predicts the problem: since the background facts of ordinary life differ from those of abstruse scientific contexts, there is no basis for expecting the same inferential schemes to work in both contexts.

Chapter 3 looks at the idea of replication of experiment. It is routinely touted in the scientific literature as the “scientific gold standard.” We find that merely a useful, but defeasible rule of thumb. It has not been given a precise enough formulation, comparable to those of the schemas of deductive logic, that would enable its mechanical application. Through a series of case studies, we find that the rule is defeasible and has been overruled in every possible combination. Successful replications (intercessionary prayer) and failures of replication (Miller experiment) have both been discarded as evidentially inert. However, on a case by case basis, warrants for the strong inferences associated with individual replications can be found in particular facts in their domains. A general principle of replication is superfluous.

Chapter 4 investigates analogy. It is a traditionally recognized argument form whose history extends back to Aristotle. However, as a review of the recent literature shows, efforts to express the form precisely as a universal rule devolve into an explosion of divisions into special cases and further qualifying clauses. Each expansion produces new problems that require further expansions and, paradoxically, carries us farther from any final formulation. This conception of analogy as an argument form is contrasted with how analogies are treated by scientists. For them, analogies are facts. This fits with a material analysis, for it allows analogies to be both facts and warrants for inductive inferences. Among these warrants, there can be no universal, formal rules. Efforts to adapt a candidate analogical rule to real examples will force a proliferation of conditions, while the rules seek a unity not present in the details of the examples. Instead, the inferences we label analogical are warranted by the facts of analogy identified by the scientists. In the examples explored in the chapter, Galileo infers analogically to mountains on the moon. His inferences are justified by the fact that the dark patches visible on the moon's surface are formed by the same processes that produce shadows on the earth. The same factual basis for inference is found in two further case studies: Reynold's analogy in transport phenomena in fluid engineering and the liquid drop model of the nucleus of an atom.

Chapter 5 takes an unflinching look at the now fashionable talk of "epistemic values" or "epistemic virtues." An early twentieth century quantum physicist who prefers the logically inconsistent old quantum theory does so, we are to suppose, because that physicist values simplicity over the competing virtue of logical consistency. The latter, however, is valued more highly by the classical physicist who then finds a different import for the same evidence. If the terms "virtue" and "value" have their usual meanings, they are ends in themselves and can be freely chosen by us. With this understanding, the physicists' inferences cease to be objective. The bearing of evidence merely reflects the physicists' freely chosen biases and prejudices. This, I maintain, is not how notions of simplicity and logical consistency are used. They are not values, but criteria, whose use is justified by their heuristic ability to lead us to the truth. They are defeasible and can be discarded when they cease to serve this end. Unless we wish to endorse an inductive skepticism by our use of tendentious language, we should stop using the misleading language of virtue and value. The term "criterion" serves better.

Chapters 6 examines the inductive criterion of simplicity in greater detail. There is no precise rule that tells us when to prefer simpler hypotheses. The later misattribution to William of Ockham, "entities must not be multiplied beyond necessity," is vacuous without specification of what counts as an entity and which are the necessities. We are bluffed into allowing its vacuity to pass because of the faux dignity of its expression in Latin. Instead appeals to parsimony in real evidential situation are abbreviated appeals to specific background facts that tell us which are the simplest cases. In curve fitting, for example, straight lines are not necessarily the simplest starting point. If we are fitting trajectories to comets, background facts tell us to start with parabolas, then ellipses and then hyperbolas. For tidal data, we start with an elaborate set of sinusoidal curves whose periods are adapted to the physical parameters of the tidal processes.

Chapters 7 probes the Akaike information criterion, which has been offered as a vindication through statistical theory of a general principle of parsimony. Closer scrutiny reveals that the criterion employs no presumption of parsimony in its derivation and that it does not entail any such general principle. Its celebrated formula merely adds a term that corrects for the overfitting of data in curve fitting problems by extra variables. We, not the statistics, illicitly interpret this narrowly applicable term as a vindication of a broader principle of parsimony. The presence of the term itself depends upon strong background assumptions, most notably that the true curve lies within the model under test. Assumptions like these are the material facts that warrant inferences that use the Akaike information criterion.

Chapters 8 addresses the popular argument form, inference to the best explanation. The hope of its proponents is that there is some feature, peculiar to explanation, that can power inductive inferences. The analysis proves unable to find such a feature. Indeed notions of explanation are so varied that instances of inferences to the best explanation may bear only superficial similarity to one another. At this superficial level, these arguments share a rudimentary common form. Real examples in science commonly begin as comparative arguments. One hypothesis is favored over another because the first entails the evidence. The competing hypothesis fails the evidence. It is either refuted deductively by the evidence or must take on a substantial evidential debt in the form of further unsupported assumptions, if it is to remain compatible with the evidence. The success of the favored hypothesis does not rest on any peculiar explanatory prowesses, but merely on its adequacy to the evidence and, more importantly, the failure of the competitor. The more fraught subsequent step of the inference must show that the favored hypothesis prevails over not just this one explicit competitor, but against all. It is often left tacit in real cases in science.

Chapter 9 seeks to reverse a decline in the literature on inference to the best explanation. This literature began rich in real examples drawn from science. The most notable is Darwin's self-conscious use of the argument form in his *Origin of Species*. Since then, proper study of scientific examples has been replaced gradually by imperfect mentions of them that often oversimplify and misinterpret them; and by prosaic illustrations drawn from everyday life. The entirety of Peter Lipton's canonical monograph, *Inference to the Best Explanation*, contains only one example from real science that is developed at length. It is Semmelweis' identification of the cause of childbed fever. The example is poorly chosen since it one of few that happens to be treated more precisely by the simple thinking of Mill's methods.

This literature is increasingly dominated by superficial examples. The best explanation for footprints in the snow is that someone walked past. This example is unlike those in science, for the human explanation of a person making distinctive marks has no serious competitors. Worse it encourages explanation by intelligent intervention. That would be an unwelcome encouragement to Darwin. He sought to overthrow intelligent creation as an explanation for biological features. My contribution is provide a somewhat more detailed exposition of eight cases in science, to which the loose pattern of inference to the best explanation can be fitted. I show in each case how some powerful, primitive notion of explanation plays no role. These examples illustrate and support the general claims made in Chapter 8 for the structure of inferences to the best explanation in real science.

With the Chapters 10 to 16, the narrative takes a different turn. The Bayesian approach presently dominates thinking about inductive inference in the philosophy of science. According to it, relations of inductive support are recoverable in some manner from probabilistic relations among proposition. I have no quarrel with the use of these probabilistic methods in domains where the background fact specifically authorize them. There are many such domains. Where I differ from the Bayesians is over their ambitions of providing a universally applicable understanding of inductive relations. It is not, contrary to the title of Jaynes' Bayesian manifesto, "The Logic of Science." It is only the logic of certain special cases. My arguments against those ambitions of universality are laid out in these chapters.

Chapter 10 is entitled "Why Not Bayes." It is a statement, not a question. I illustrate how background conditions can lead us to non-probabilistic representations of evidential relations using the extreme illustration of completely neutral evidence. For this case, application of simple invariances lead to a highly non-additive representation of inductive support. It is quite contrary to the additivity of a probability measure. I argue that even the contrivances of the new literature in "imprecise probability" can sometimes fail to do justice to it.

Bayesian analysis is distinctive in that, laudably, it has taken seriously the burden of proving the uniqueness of its probabilistic representations. This chapter argues that all these efforts must fail since they all have the same structure. Whether they are Dutch book arguments or employ representation theorems, they proceed from some set of assumptions and then *deduce* that the targeted beliefs or relations of inductive support must conform to the probability calculus. This last conclusion is a contingent proposition. It follows that it can only be deduced from assumptions that are at least as strong as it logically. Hence, necessarily, the assumption of probabilities must be hidden within the starting assumptions. The proofs are not demonstrations of the necessity of probabilities, but merely a restatement of a preference encoded in its premises. Once one realizes this, it merely becomes a mechanical exercise to identify and expose the hidden assumptions. I carry out the exercise for Dutch book arguments and representation theorems and note that all similar arguments will fail in the same way.

Chapter 11 contains an extended example of this last exercise. The scoring rule or "accuracy" based vindication of probabilism is based on a dominance theorem. If our credences are not probabilistic, then the theorem tells us we can always improve the accuracy of our credences, no matter what the true situation may be, merely by shifting our credences to a probability. The chapter shows that the theorem is sensitively dependent on the particular scoring rule used to measure the inaccuracy of credences. It develops a family of scoring rules such that any desired deviation from additivity in the credences can be secured merely by choosing the requisite rule from the family. Then a variant theorem shows the dominance of credences with the specified deviation from additivity. The literature in accuracy-based vindications has sought to parry such possibilities by seeking further reasons for why only those rules that deliver probabilities are admissible. These efforts cannot succeed since they still seek to derive probabilities deductively from further assumptions. I continue the exercise of displaying how these further assumptions still have hidden within them the presumption of probabilities.

Chapter 12 addresses a more general problem facing all efforts to devise a mathematical calculus for strengths inductive support. Applications of Bayes' theorem require specification of prior probabilities. They make a difference to the resulting posterior probabilities. Since they must be determined by factors external to this application of Bayes' theorem, it follows that this specific computation is not inductively self-contained. One might hope to eliminate this dependence on external considerations by a suitable expansion of the scope of the application of Bayes' theorem. The present prior probabilities would then be recovered as posterior probabilities of antecedent applications of Bayes' theorem. Continued expansion might, we hope, eventually eliminate the intrusion of external considerations. It is well-known that these hopes fail. No matter how large the scope of the application, one is never freed from the need to use external consideration to fix prior probabilities.

It turns out that this inductive incompleteness of the Bayesian system is not a failure unique to the Bayesian system. Rather, it is an instance of a broader incompleteness that afflicts all candidate calculi of inductive inference. That is, a theorem demonstrated elsewhere shows that this incompleteness must arise in all such calculi that conform with weak and broadly acceptable conditions. This chapter does not develop the theorem in all its mathematical details, but presents its core ideas and some illustrations of it. The theorem gives a precise instantiation of the more nebulous slogan, "there are no universal rules of inductive inference." It shows that there are no inductively complete calculi of inductive inference.

The remaining Chapters 13 to 16 display further situations in which the background facts warrant formal treatments of inductive support that are not probabilistic. They illustrate the locality of inductive inference. In each case, we must first find the facts prevailing in some domain and then read from those facts the particular logic that would apply to it.

Chapter 13 considers an infinite lottery machine that chooses without favor among a countable infinity of outcomes, labeled 1, 2, 3, 4, ... The condition that the lottery machine chooses without favor is expressed as an invariance, "label independent." According to it, the support accrued to any individual outcome, or set of outcomes, remains the same no matter how we may permute the labels. This independence exercises a profound restriction on the formal behavior of strengths of support. For example, all infinite sets of outcomes whose complements are also infinite must accrue the same support. This sector of the logic is highly non-additive. A corollary is that the relative frequency of even numbered outcomes does not stabilize towards one half in many, repeated drawings. Rather, all relative frequencies continue to accrue equal support. The factual conditions characteristic of the infinite lottery machine turn out to arise in a particular problem in recent inflationary cosmology. The infinite lottery machine logic is the applicable logic.

Chapter 14 undertakes the same exercise for an uncountably infinite outcome set, such as the continuum-sized set of outcomes formed by the real numbers between zero and one. One might think that choosing without favor among outcomes in this set is easily achieved probabilistically by a uniform probability distribution. That is a misleading impression since, by foundational design, such a probability distribution neglects to assign probabilities to very many subsets of outcomes of the space. If we require a representation that covers all subsets, we arrive at a logic similar to that of the infinite lottery machine logic, but with more sectors. The chapter then considers successive restrictions that would move the logic towards a probabilistic logic. With each restriction we find a variant of the non-probabilistic inductive logic warranted. One application of these intermediate logics is the continuous creation of matter in the steady state cosmology of Bondi, Gold and Hoyle. The most interesting cases technically arise with paradoxical decompositions of measure spaces. They show the existence of outcome sets not measurable by additive measures such as a probability measure. To make their character more concrete, the chapter develops nonmeasurable sets derived from coin tosses. It turns out that a variant, but weak inductive logic—an "ultrafilter logic"—applies to these sets.

Chapter 15 investigates the inductive logic warranted in two sorts of indeterministic physical systems. The first are those whose temporal behavior is indeterministic. They are quiescent for an arbitrary time and then, without any specific triggering event, spontaneously move. The chapter develops the especially simple example of the infinite domino cascade, which is new in the literature. The second type of indeterministic system is those in which specification of one part of the system fails to fix the remainder. Fixing the mass distribution in Newtonian cosmology fails to fix the gravitational potential. It is then shown that no probability measure can represent the indeterminacy. The infinite dimensionality of the space of Newtonian potentials presents especially intractable problems for additive measures. Instead, it is shown that the background facts of the systems realize the invariance that led to the completely neutral support elaborated in Chapter 10. This is the logic applicable to these indeterministic systems.

The alternative inductive logics explored so far all tend to be simpler in their structures than the additive measures of probability theory. Chapter 16 shows that this need not be so. The system considered is the spin of electrons in quantum theory. While probabilities arise in the process of quantum measurement, they turn out not to be the structure representing inductive support that is warranted by the physical facts of quantum theory. Rather that structure is the density operator that also represents states in quantum theory. The chapter explains what these operators are, how they come about and how they represent inductive support. The development is written at a level that presumes no special knowledge of quantum theory, but assumes a little comfort with abstract mathematics. We learn from the example that background facts in some domains can warrant an inductive logic of some complexity that is quite different in its structure from a probabilistic logic.

5. "A" or "The"?

Finally a note on terminology: is it *a* material theory of induction or *the* material theory of induction? I use both expressions. The first refers to the general idea of finding the warrants for inductive inferences in background facts. There is no presumption in this usage of a particular way of proceeding beyond just the general idea. The second expression—*the* material theory of induction—refers to the particular instantiation of the general idea found in this book and my relevant papers.

References

Norton, John D. (2005) "A Little Survey of Induction," in P. Achinstein, ed., *Scientific Evidence: Philosophical Theories and Applications*. Johns Hopkins University Press. pp. 9-34.