Author's responses

John D. Norton

Department of History and Philosophy of Science, University of Pittsburgh, USA

1. Introduction

To see an issue of this journal dedicated to the material theory of induction is not merely satisfying. It is quite thrilling. It tells me that there is considerable interest in the material approach to inductive inference and that there is much more to come. I offer my thanks to each author who contributed to this volume for their interest in the topic and for the effort needed to prepare their papers. My special thanks, however, are reserved for the volume's editors, Wendy Parker and Elay Shech. Assembling a volume like this may seem like a minor chore to those who have never tried it. The reality is that the work is trying, exhausting and endless. They persevered and they surely must be pleased with the result.

In reflecting on the editors' invitation to respond to the papers in this issue, it became clear that no unified response is possible. Each paper picks up a different aspect of *The Material Theory of Induction*. Each requires a separate response. A unified response would be reduced to generalities that fail to connect properly with any. Those individual responses are provided below. Many of the papers have mounted challenges to positions I defend. This is to be expected. My volume disputes wisdoms across the board that have become established with greater comfort than they deserve. The challenges in this issue are sometimes spirited. I have not shied away from responding with comparable vigor. These papers are not exercises in perfunctory adulation. They are a stress test of the material theory of induction. My sense is that the theory passes the test and I explain why I think so in my responses.

2. Alan Baker, "Schemas for induction"

One of my critiques of the formal approach to inductive inference is that all inductive schemas face counterexamples. The natural reaction to these counterexamples has been to embellish the schemas to preclude contradiction with the counterexamples. What results is a cycle of woe, a death spiral. Each embellished schema faces new counterexamples. Efforts to further embellish the schema to accommodate them attract new counterexamples; and so on, apparently without end. With each iteration of the cycle, the scope of the schema becomes smaller while the formulation of the schema becomes longer and more convoluted.

My diagnosis is that the embellishments can make momentary gains since they are enriching the schemas with background facts. As we add more, these facts play a greater role in the warrant supplied by the schema. The material theory simply takes this process to its endpoint. All the authorization of an inductive inference is carried out by background facts.

Alan Baker has proposed an ingenious escape from the death spiral. Do not succumb to the natural urge to embellish the schemas. Simply accept them for what they are, fallible. That is, he urges that the counterexample ridden schema provide us with a "defeasible license" for an inference. We should accept what they license as a default until we have reasons to do otherwise. Those reasons would be the sorts of counterexamples that I have displayed as troubling the schema.

This reverses my view that no schema can claim the status of a default. We have a positive obligation in each case to establish the appropriateness of the schema to the case at hand. That appropriateness is established by displaying a suitable background fact in the context.

Two things concern me about Baker's proposal. First, adopting a schema as a default is only viable if the schema is *mostly* correct in its applications. To use Baker's example, we can say that *ceteris paribus* lightning strikes the tallest object since that is what mostly happens. However, the sorts of schemas on offer are unlike this. They mostly fail. Take Baker's "EIS" schema:

All observed F's have been G.

Therefore, the next observed F will be G.

If we are truly free to insert anything into the slots of F and G, then this schema will almost always deliver a false result.

This failure might not be apparent from Baker's narrative. He gives as an instance the familiar:

All observed emeralds have been green.

Therefore, the next observed emerald will be green.

This example is exceptional. The F's and G's have been chosen very carefully, under the guidance of background facts, so as to ensure that the inference works. This appears natural since no details are given of the observations. If we had a fuller account of just what the observations were, we would find a vast array of F's and G's for which the inference would fail.

For this reason I prefer the example I give in Chapter 1 of *The Material Theory of Induction*. Curie observed that a tiny sample of radium chloride in her laboratory is crystallographically like barium chloride. She infered all are so. This one fact could be redescribed in very many ways, using very many different "F's" and "G's." Is the observation of a tiny thing in her laboratory that has spiky crystals: or a crystal that is at temperature 20C; or a crystal in Paris; or a crystal prepared by Curie? I go to some lengths to indicate just how careful Curie had to be in picking her predicate G. The inference would fail for almost every other possible candidate other than the carefully chosen "crystallographically like barium chloride."

My second concern with Baker's proposal is that he gives us no good reason to accept any inductive schema as a default. In many cases of deductive inference schemas, those reasons are provided by the meaning of the terms. Take:

All F's are G.

Therefore, some F's are G.

That this schema generates valid inferences follows from the meaning of "all" and "some." Nothing more is needed. However what basis do we have for taking the inductive version of this schema as a default?

Some F's are G.

Therefore, all F's are G.

The inferences this schema authorizes amplify the factual content of the premises. Why should we expect this amplification to succeed? Are we to assume that once an F has appeared with a G, the world is such that it will generally continue to do so? If we can substitute anything for F and for G, it is surely not the case. I argued in Chapter 2 of the *Material Theory of Induction* that this amplification succeeds only when we implement the schema in an environment hospitable to the amplification for these specific F and G. That the environment is hospitable, I argued further, is a factual matter that may be true or false. That it is true, when it is, provides the material warrant for the inference.

These two concerns taken together lead me to prefer my material diagnosis of the prevalence of counterexamples to formal schemas of inductive inference.

3. Paul Bartha's, "Norton's material theory of analogy"

Paul Bartha has written a careful and thorough appraisal of the material analysis of analogical inference. I agree with almost everything he says. He disavows a fully formal account of analogical inference and suggests that we seek one that has both formal and material elements. Such an account seems quite possible to me. Analogical inferences often share formal similarities that allow them to be grouped together. This fact has encouraged efforts to find formal accounts of analogical inference. Such an intermediate account could be quite useful in practical investigations in the ways that Bartha has outlined. Bartha also correctly

points out that scientists, engaged in analogical reasoning, might not always start with a fact of analogy. That source and target system are governed by some particular fact of analogy may only become explicit at the end of the analysis.

Much of this, however, goes beyond what I seek to establish in the analysis of analogical inference. I do not seek to reconstruct the procedures used by scientists in their efforts to discover good analogical inferences. That project is one of far greater ambitions than mine. While I hope that the material perspective will be useful practically to those seeking analogical inferences, an account of that practice requires much more extensive analysis such as Bartha has given in his writings.

My concern is limited to justification or warrant. What makes some particular analogical inference a good one? My answer is that the chain of justification terminates materially, in a fact of analogy. The inference is only good in so far as that fact is a truth. The justificatory chain may pass through intermediate formal schemas of limited scope. But those schemas, as applied to the case at hand, will in turn be justified by the pertinent fact of analogy.

With these limited ambitions recalled, I can now explain why Bartha's specific criticism is untroubling. His first and main objection is that, he says, the material account has limited scope. He lists (Section 5.1) three scenarios in which he asserts that the material account fails. They are:

(a) No clear fact of analogy.
(b) No independent support for the fact of analogy.
(c) The fact of analogy is too vague to support any inference.

The concern in (a) is that "There may be no specific identification of any *fact of analogy*." If there is no warranting fact of analogy, then the candidate analogical inference is an inductive fallacy. That is the end of it. If there is a suitable fact of analogy, there is no need for the scientist to identify it. The analogical inference is warranted, even if the scientist may not be able to articulate the warrant or even know it. This phenomenon is familiar in deductive scenarios. Many scientists are not schooled in deductive logic and cannot name many or even any deductive schema. Yet they are competent in deductive inference. Matters can be more abstruse. How many physicists know that they are relying on the axiom of choice when they attribute a basis to some large vector space? Or that they are presuming the axiom of countable additivity when they sum infinitely many probabilities? Imperialist Bayesians will insist that all the scientists' evidential reasoning is justified, unbeknown to them, by the axioms of the probability calculus.

The concern with (b) is "Even if a *fact of analogy* is identified, there may be no independent inductive argument offered in support, and there may be none available." The fact of analogy must be a truth if the analogical inference is to be warranted. If it is not so, then the candidate analogical inference is an inductive fallacy. Bartha provides a good example. The discrete lines in the emission spectrum of hydrogen were puzzling. In 1871, Stoney (1871, p. 296) suggested that the discrete frequencies arose from oscillatory molecular motions analogous to the specific harmonics produced by a violin string. The fact of analogy presumed is not true. The discreteness of the lines is quantum mechanical in origin. They are associated with the quantum jumps of Bohr's 1913 theory of the atom and the mode of their production was eventually recovered from quantum electrodynamics. Hence, I do agree with Bartha's conclusion: "So, on Norton's account, there is no warrant for the analogical inference." That is just as it should be.

Finally, matters are the same with concern (c). If the fact of analogy is too vague to support the inference, then it is not warranted. Once again, that is just as it should be.

Bartha's main criticism derived from his misunderstanding of the material theory as an account of the procedures scientists use in arriving at analogical inferences. His second criticism asserts a limitation on the range of analogical inferences that the material account can judge to be warranted. As far as I can see, it amounts to this (Section 5.2): "the

analysis simply doesn't work if the fact of analogy is not patent, not independently justifiable, or does not function as an intermediate step in the argument." I find the analysis to work well in all these cases. The fact of analogy can serve as a warrant even if it is not patent or explicitly present in the argumentation offered by the scientist. If the fact of analogy is not justifiable, then the candidate analogical inference is an inductive fallacy. There is no other way to warrant it. For reasons already given in my chapter, it is hopeless to seek a warrant in some formal schema or even in what Bartha calls a "quasi-formal" theory.

4. Pat Corvini, "What induction is (and what it should not be): a concepts-centric perspective on Norton's radium chloride example"

Pat Corvini has taken rather sharp exception to my treatment of Curie's inductive inference in Chapter 1 of the *Material Theory of Induction*. The basis of our disagreement lies in our different conceptions of the nature of inductive inference.

For Corvini, what makes an inference inductive, at least in these cases, is that it infers from particulars to a generality. What is essential to the cogency of the inference is the use of appropriate concepts. Inductive risk may be present to a greater or lesser degree. However, it is a nuisance distraction. She writes: "I thus believe that focusing on 'risk' as an essential marker of '*induction*'—as Norton's analysis of this example would have us do—is a mistake."

For me, as Chapter 1 makes clear, inductive inferences are defined as those that arrive at conclusions deductively stronger than their premises. Inductive risk is not an incidental feature. It is constitutive of inductive inference. This broad conception is called for by the wide range of inferences and inferential practices covered in the sixteen chapters of the book.

What may be lost in this emphasis on our differences is that Corvini and I largely agree on the example. Much of my chapter is devoted to making just this point. Inductive generalizations like Curie's can only work when we have found what are, in later terminology, the "projectible" properties. I went to some pains to stress how difficult it was to find the appropriate projectible properties in crystallography and to criticize approaches that neglected this difficulty.

How can this episode in crystallography be used by each of us to illustrate our views on inductive inference? For Corvini, the key is the formation of the appropriate concepts. It is a daunting task to give a proper account of the establishment of these concepts. By Curie's time, that was the work of over a century and well beyond the scope of an introductory illustration in an introductory chapter. Instead I chose a simple inductive inference that came after the hard work of establishing these projectible properties was over. Curie and the other radiochemists of the early twentieth century found that samples of radium chloride have particular crystallographic properties. They then infer inductively that all do.

The inference is a small step in the larger crystallographic project. That is what makes it tractable in an introductory illustration. It means that we can have full control of the inductive risk taken. The example illustrates the key idea of the material theory of induction: how background facts warrant an inductive inference.

While this example serves my purposes, it does not serve to illustrate Corvini's approach to induction, since it comes after all the hard conceptual work has been completed. So I understand her discomfort with my choice of the example. While it illustrates my point, it does not illustrate hers. Unfortunately, Corvini has presented my choice of example in rather dark tones. She calls my treatment of the example not merely mistaken, but, repeatedly, one that "misrepresents" the ampliation. The term has connotations of dishonesty and deception, and, in the legal context, fraud. Similarly, my historiography is dubious, harboring "historical inaccuracies" and "a frustratingly complicated tangle of fiction and fact."

These are serious accusations and I took them seriously. However,

after considering her many objections, I do not find that any of them substantially alter my analysis or the conclusions I have drawn. They do call for one minor amendment.

To recapitulate, Curie and other radiochemists at the start of the 20th century had prepared only few samples of radium chloride. They noted the uniformity in the crystalline forms and described them using the vocabulary of mineralogy of the time. In reporting their observations, they made a generalization. I reported the generalization as an unqualified "All samples …." After Corvini's critique, I can see that they likely limited the scope of the "all" tacitly. I would change the generalization to "All samples prepared under comparable conditions …" This is a generalization from the few actual cases at hand to all cases, now restricted by a "comparable conditions" clause.

My basic point remains. In making the generalization, the radiochemists exposed themselves to inductive risk. Given the accepted mineralogy of the time, if radium chloride was to form crystals at all, they would have to be in one of the identified crystallographic families. Under it, the only thing that can go wrong in the generalization is that the crystals do not form in the monoclinic family in which barium chloride also crystallizes. That is the inductive risk of polymorphism that I identify.

It is a real risk. The tacit restriction to "comparable conditions" is a vague gesture at something of great complexity. We must answer many questions to specify it. What is the temperature of the solution in which the crystals form? What is the rate of evaporation? What is the composition and texture of the walls of the containing vessel? What is the mode of heating? Is it uniform or localized? What are the possible impurities? Are there crystallographically anomalous seed crystals already present? In coming to their general conclusion, Curie and the other radiochemists of the time were taking the risk that changes in one or more of these conditions might produce different results. It seems fair to say that the risk is small, but it is there.

The basic moral I draw from the illustration remains. The inductive inference Curie and the other radiochemists of her time made was warranted not by a formal schema, but by a background fact. I have called it the "weakened Häuy's principle" and Corvini may want to dispute the name since Häuy himself worked nearly a century before Curie.[1] However its content does summarize the background facts Curie needed to warrant the generalization.

5. Kevin Davey, "Inference to the best explanation and Norton's material theory of induction."

The goal of Kevin Davey's paper is to offer a friendly amendment to the material theory of induction. The context dependence of inductive inference is to be retained, but the locality of rules of inductive inference is to be replaced by a universally applicable schema of inference to the best explanation ("IBE").

A major part of Davey's text criticizes a reconstruction of IBE arguments in Chapters 8 and 9 of *The Material Theory of Induction*. I agree with the criticism. To explain why, we need to recall the structure of the main argument in those chapters. The schema of IBE depends on the idea that hypotheses do not accrue support merely by accommodating evidence. They have to explain it. There is, we are to suppose, some especially potent inductive power in explanation, not present in mere accommodation. When I began working on IBE, I was unable to identify that extra power, intrinsic to explanation, or even to come to a clear idea of the notion of explanation at issue. Since a large repertoire of instances of IBE in science is available, I decided to proceed inductively. I collected many of these instances and tried to discern how they worked. The

---

[1] Tutton (1910, p. 4) attributes to Häuy a "great truth" that he calls a "principle" and is given by him in italics: "to every specific substance of definite chemical composition capable of existing in the solid condition there appertains a crystalline form peculiar to and characteristic of that substance."

results are reported in Chapters 8 and 9. I identified a common structure in which there was no distinctive notion of explanation. Crucially this commonality persists only at a broad and superficial level. An appraisal of the inductive strength of each instance requires examination of the details of each and the particular background facts warranting them.

The common structure consists of two steps, detailed in the chapters and Davey's text. The second step is that the better of competing hypotheses is to be taken as the best, absolutely, and inferred. This Davey calls the "better-is-best principle." Davey raises significant objections to the principle. He asks, in good material fashion, what material facts could justify it. He finds none adequate. He asks: "Why does Norton think that his better-is-best principle is justified?"

There is an easy answer to this question: I do not think it is justified. I agree with Davey's criticism. Indeed, I do not think that I even identified a principle. All my narrative gave was a summary of the commonalities among standard IBE's in science. The mere fact of carrying this superficial similarity is not enough to warrant an inference. Materially understood, each instance identified has to be warranted by suitable background facts. In this case by case analysis, the second step, I expected, would prove to be the greatest weakness of the arguments reconstructed. My summary judgment, quoted only partially in Davey's text, (Ch.8, §1) was:

> The second step is more fraught. We are to suppose that better is best; and that best is good enough to warrant commitment. Preference becomes commitment. The step is commonly grounded in a presumption that no other theory can do better than those explicitly considered. That presumption is so hard to justify that this second step is often left tacit and sometimes even omitted completely. For the step commonly relies merely on our human imaginative powers to sustain the conclusion that there is no better account just beyond our horizon. Kyle Stanford [reference] has effectively and powerfully described this problem of "unconceived alternatives."

Then, in Chapter 9 (§3), where I summarize the instances in a table, the column summarizing that step has entries like "tacit," "not taken" and "complicated."

These are not ringing endorsements of Davey's principle. They convey my repeated concern about the weakness of this step in each case. However, I forgo a definite, blanket statement against the step since the material theory of induction enjoins me to treat each case on its individual merits.

The pertinent conclusions of the two chapters are that (i) the standard examples of IBE in science do not employ any distinctive notion of explanation with peculiar inductive powers; and (ii) these standard examples share a common weakness in so far as they try to carry out a better-to-best inference. As a result, I find it hard to share Davey's enthusiasm for IBE. The enduring and apparently irremediable difficulty is that the notion of explanation employed remains obscure, as does the origin of its inductive powers. The case for this difficulty is developed at length in Chapter 8, so I do not need to rehearse it again here.

Davey dismisses this concern that, as he puts it, "IBE is so vaguely defined as to lack real substantive content." He continues that it is "unfair, as at least nowadays there are many reasonably precise formulations of IBE." This appraisal is supported by three citations. None, however, give a precise account of explanation, adequate to sustaining IBE. One (Douven [2018] in Davey's references) writes:

> IBE is best thought of as a slogan that can be fleshed out in different ways, where different fleshings-out may have different merits and drawbacks, depending on the context of usage … it is difficult to find a statement of IBE that is more specific than the slogan-like characterization …

Instead of providing some reasonably precise universal formulation, this appraisal aligns better with the contextuality of the material analysis. The other two citations are to Bayesian analyses of explanation.

Bayesian analyses face many difficulties. One is already sufficient to rule them out for present purposes. Bayesian analyses cannot have universal application since only a subset of relations of inductive support or of credences are probabilistic. Chapters 9–16 of *the Material Theory of Induction* are devoted to demonstrating this limitation. Curiously, Davey agrees with my hesitation over Bayesian analyses. His footnote reads "… I actually do not think that developing IBE within the context of a probabilistic framework is the right approach. The problem is that IBE seems to be applicable in cases in which probabilistic machinery is not."

Davey seeks to excuse what seems to be a fatal imprecision by suggesting that we are asking too much. "Norton does not tell us anything positive about what the 'licensing' relation is …," he writes. Here I see no lacuna. Material facts license inferences simply through the meaning of the propositions that express them. No lofty and general philosopher's account is needed. It is the same with deductive logics. The proposition "If *A* then *B*." authorizes us to infer from *A* to *B* simply in virtue of the meaning of "if … then …" Correspondingly the material fact from Chapter 1, "Generally, each crystalline substance has a single characteristic crystallographic form." warrants fallible generalizations concerning the properties of samples of crystalline substances in virtue of the meaning of the proposition. If, however, there is an imprecision here, then IBE faces compounded imprecisions: first, in the notion of explanation itself; and then in accounting for how explanation licenses inferences.

Overall, far from providing a universally applicable inductive schema, my concern is whether IBE can, even in individual cases, go beyond assertions about explanation in vague, general terms and give us a serviceable, local schema.

Finally, a misunderstanding in Davey's text needs to be corrected. It labors over the problem of whether the material theory of induction is (or assumes) an internalist or an externalist epistemology; and settles on an externalist epistemology. The correct answer is "neither." The distinction between internalist and externalist epistemologies applies to accounts of the mode of justification of beliefs held by some agent. The material theory of induction concerns relations of inductive support among propositions, independently of whether these propositions are held as beliefs by some agent. Here the theory proceeds as would an ordinary logic textbook. There modus ponens is a valid form and affirming the consequent is not, independently of whether the propositions they relate are held as beliefs. Perhaps assertions of this independence in Chapter 1 of *the Material Theory of Induction* have been misread as endorsements of an externalist epistemology in which the justification of an agent's beliefs are inaccessible to the agent? Of course, some agent may use inductive inferences supplied by the material theory of induction to justify belief in some proposition from beliefs in others. The details of the resulting system of justifications, however, are left by the material theory to epistemologies of beliefs.

6. Job de Grefte, "Epistemic benefits of the material theory of induction"

According to the problem of induction, attempts to justify a rule of inductive inference are either circular or trigger an infinite regress. The material theory of induction, I maintain, dissolves the problem, since it has no universal rules of inference in need of justification. Attempts to set up an analogous regress problem in the justification of material facts fail, since they neglect the non-hierarchical structure of relations of inductive support. It was encouraging to see that Job de Grefte accepts this dissolution in his paper. Yet, he is unimpressed. He argues that the problem of induction has already been solved in externalist epistemologies and that the material theory of induction offers no epistemic advantages over the externalist solution.

The primary goal of the material theory was to determine which are the good inductive inferences and what makes them so. The material dissolution of the problem of induction was an unexpected, secondary benefit. Nonetheless, de Grefte's appraisal underestimates the value of

the dissolution. His analysis rests, as far as I can see, on three claims:

1. The problem of induction is most fundamentally a problem in the epistemology of belief, where its solution is to be sought.
2. Externalist epistemologies have solved the problem in that context.
3. A material theory of induction requires an externalist epistemology.

Here I will explain why I doubt each of these claims.

First, we need a terminological clarification: De Grefte is concerned that my use of the term "inference" is ambiguous, when it is not. In his epistemological literature, "inference" refers to the passage from one belief to another and "implication" refers to logical relations among propositions. As I explain in the introduction to Chapter 1 of *The Material Theory of Induction*, my use of "inference" as a purely logical relation among propositions follows standard logic texts of the nineteenth and twentieth centuries. Here I will continue to resist his annexation of the term and will continue to use "inference" in this long-established logical sense. This maintains a continuity of terminology with the logical tradition in which I work. To shift its meaning, as de Grefte and others in the epistemology of belief seek, is to invite confusion. Someone in one tradition can then readily misunderstand the claims of someone in the other. Might just such a confusion be the origin of de Grefte's complaints about my treatment of the problem of induction?

On 1.: The version of the problem of induction addressed by the material theory resides fully within the study of inductive logic and concerns its rules of inference. The rule of enumerative induction, the problem tells us, cannot be justified by noting that the rule has always succeeded in the past, for that is circular. This version of the problem and the relations explored by inductive logic are independent of our beliefs; and that is an important fact. Whether the fossil record provides strong inductive support for the theory of evolution is independent of our thoughts and beliefs. Hence the analysis of this version of the problem should proceed independently of our beliefs; and that is just what the material dissolution does.

De Grefte's proposal is that the material dissolution is superfluous since an analogous version of the problem has been solved by externalists in the epistemology of belief. His claim is quite unambiguous and depends, as far as I can see, on a simple equivocation. In a section entitled "4. The problem of induction is an epistemological problem," he writes:

> My aim is to problematise Norton's dissolution of the problem of induction. In this section, I will argue that this is an epistemological problem. As we saw in the previous section, the material theory is best understood as a theory of logic. That means that in order to bring the material theory to bear on the problem of induction, we have to specify first how a theory of logic can have epistemological implications.

> What follows is the suggestion that the material dissolution is superfluous since externalists have already solved the epistemological problem.

Of course, I agree that we need to see how the material theory of induction bears on beliefs if it is to help us solve a problem concerning beliefs. What I do not agree with is the equivocation when de Grefte's shifts the meaning of the term "problem of induction" in the phrase "Norton's dissolution of the problem of induction." The target of my dissolution is the logical problem of induction. De Grefte's treats my dissolution as an attempt to solve a different problem in epistemology and is then dissatisfied with the result.

The only defense I can see is if de Grefte has somehow come to believe that the logical problem derives from an analogous epistemological problem that bears the same name. Then solving the epistemological problem might well be needed to solve the logical problem. In an earlier draft, on the basis of some remarks in de Grefte's text, I wrote that de Grefte had argued that logical relations among propositions derive

from our beliefs about them. In correspondence on the draft, he assured me that he does not hold this indefensible belief. This leaves me unable to see what basis de Grefte has for "problematising," as he put it, my logical dissolution of a problem in inductive logic. Indeed, I am unable to square it with his later remark "As always, we must be careful not to confuse logic and epistemology."

On 2.: Since externalism in the epistemology of belief is not a view I know well, my comment is brief. To know that our beliefs are justified in a reliability externalism requires that we know that our processes of belief formation are reliable. We cannot know that they will continue to be reliable on the evidence of their reliability in the past. Otherwise we commit precisely the circular reasoning of the problem of induction. That they might be reliable, without us being able to show it, is irrelevant. We have no assurance of it. In this regard, externalists are no better off than those internalists who proceed only with the hope that future applications of their methods are truth conducive but cannot justify it.

On 3.: The material theory of induction has no special connection to either internalist or externalist epistemologies. It is a resource that both can use. As de Grefte notes, it is comforting to externalists that someone can carry out a good inductive inference without knowing the warranting fact, which is external in that sense. However, internalists can respond that the warranting fact is accessible and that the activity of identifying such facts is routine. Much of the chapters of *The Material Theory of Induction* are devoted to this activity.

De Grefte argues that internalists can accommodate a formal account of inductive inference, but not a material account. The difference, he argues, is that a formal theory concerns the forms of propositions and those forms are available in mental content. However, the forms of the propositions are not enough to establish the validity of the argument. The forms can fit with many schemas, some licit, some not. To be assured of validity, internalists also need to know which schemas are licit. For example, from the premise that, for any prime, we can always find a greater one,[2] Euclid infers that there is an infinity of primes. When we accept that inference, we accept that some version of the deductive schema of mathematical induction is licit, even though strict intuitionists do not allow it. That requires further knowledge, just as the material theory requires us to learn further facts to be assured that some inductive inference is good.

7. John Earman, "Quantum sidelights on the material theory of induction"

It is our good fortune that this collection has been able to include Professor Earman's reactions to the material theory of induction. He is a scholar for whom I have great admiration and from whom I have learned very much, even while being his colleague and occasional co-author. There is much to learn from his article and nothing in it to dispute. He is quite right that there are many other cases of indeterminism arising in our physical theories. He can write this with some authority. His 1986 *Primer on Determinism* has defined the modern study of determinism in physics.

Professor Earman's analysis concludes with a mild rebuke: "… The Material Theory of Induction undervalues the Bayesian account of inductive inference." The remark is underwritten by a detailed study of how Bayesian ideas can be adapted to quantum theory. It is a fair rebuke, for the volume is heavily laden with complaints about Bayesian analysis.

Here I should emphasize that I fully realize and respect the immense utility of probabilistic analysis in many contexts. Modern statistical analysis without it would be unthinkable. That utility is not my target. Rather my target is the idea that Bayesianism provides the universally applicable account of relations of inductive support. This idea has

---

[2] If p is prime, then $p!+1$ is either a greater prime or, if not, divisible by another prime greater than p.

become a tacit default for many who work in the philosophy of science. It encourages them to offer Bayesian analyses of many topics in philosophy of science.

To counter this presumption that Bayesian analysis is a universally applicable default, the volume goes to some pains to identify circumstances in which Bayesian analysis is not appropriate. It also seeks more general arguments against that universality. It is inevitable that the overall impression is one of great negativity. That negativity must be understood as directed narrowly at the claim of universality. I have no doubt that probabilistic analysis can be of great value in the right contexts; and that there are many of these contexts.

I urge a reversal of approaches. In each context, we cannot presume that Bayesian analysis is the automatic default. Rather, in each case, we have an obligation to display positively why Bayesian analysis is warranted. We should seek that warrant in the facts prevailing in the context. Such is the case in Professor Earman's example. The probabilities derive from the basic facts of quantum theory. In other cases, however, the background facts will not warrant probabilities. We commit inductive fallacies if we nonetheless persist in their use. A striking example is the "inductive disjunctive fallacy" described in Norton (2008, p. 509).

8. Benjamin Genta, "How to think about analogical inferences: a reply to Norton"

Benjamin Genta's contribution to this volume addresses the application of the material theory of induction to analogical inference. It contains a catalog of what he labels "worries" and a positive proposal to guide further research.

While I can offer no guarantee that my analysis of analogical inference is without flaws, I was relieved to find that Genta's worries have not revealed any. Since his catalog is rather long, I cannot address them all. Rather I will address the three categories of worries listed in the three subsections of his Section 5.

The first worry in Section 5.1, "Ambiguous Facts of Analogy," is that the material approach does not define precisely what a fact of analogy is. The material theory does as much as can be done in characterizing them as asserting some similarity between the source and target system. That description is as precise as has been routine in the literature and has been sufficient to support a history of analogical inference extending over millennia. The more serious worry is that one inductive inference might be warranted by multiple different facts of analogy. There are, he believes, an infinity of possibilities. Yet, he asserts, the material theory of induction gives us no means to discern which is the right one. This is not so. According to the material theory of induction, the warranting fact must be a fact, that is a truth, else the inference warranted is an inductive fallacy. Thus the infinitely many candidates likely contain infinitely many falsehoods. The material theory gives ample means to discern the right one. For determining which of many candidate propositions are factual is a task for inductive inference through further evidential exploration.

The worry in Section 5.2 "the Material Theory and Normativity," is that the material theory gives us no guidance in choosing good facts of analogy and, further, "does not allow us to differentiate between good and bad." The material theory gives scientists good advice on how to proceed: seek truths. The good warranting facts will be true; and the bad ones will be false. The guiding principle is that the more you know, the more you can infer. So if you are interested in analogical inferences connecting two domains, the material theory recommends learning as much factually about the similarities between the two domains as possible. Seek the governing fact of analogy in what is learned about those similarities. Those efforts will be rewarded better than efforts devoted to pondering the general formal structure of similarity and dissimilarity.

Section 5.3 "A Possible Counterexample," offers a case of an analogical inference that, purportedly, is not warranted by a fact of analogy. The analogy is between a general laying siege to a city and a doctor treating cancer tumors with radiation. The difficulty with the example is that it is a very weak analogical inference and should be judged so by any account, material or formal. The formal theory will judge the disanalogies stronger than the analogies and discount it. A material analysis will find only a thin fact of analogy that merely notes that the two cases are superficially alike: an agent that can be divided (radiation, armies) acts on its object (tumor, city). This fact warrants little. The analogy could be strengthened if we imagine that the general wants to minimize civilian casualties; and the doctor wants to minimize harm to non-cancerous tissue. Then the critical fact of analogy would be that the collateral damage of the agent grows faster in magnitude than the magnitude of the agent. If that turns out to be the case, the collateral damage is minimized if the agents act in divided portions. In sum, the example is not so much a counterexample as a bad example.

Genta's paper concludes with a positive proposal for a "Guiding Principle of Analogical Inference." The hope expressed in the conclusion of the paper is that this principle will help convince us that "future studies of analogical inference will benefit from additional formality …" I take a dimmer view of this urge to formalism. The principle falls far short of what a formal analysis should provide. It depends essentially on a formally unexplicated notion of similarity both within each of the source and target and between them; and a formally unexplicated notion of the weights assigned to them.

According to the material theory of induction, there are no universally applicable schemas for inductive inference. The familiar schema, such an analogy or enumerative induction, are only possible because there is a loose similarity among the inferences they group together. The similarity is only superficial since the instances grouped together will vary in important details as we move from domain to domain. As a result, the schemas work near enough, but no better. The prediction of the material theory, then, is that efforts to adapt the schemas more precisely to the range of instances will result in an ongoing explosion of extra clauses and conditions, each tailored to accommodate some anomalous instance. A major theme of the chapter on analogy in *The Material Theory of Induction* is to display this explosive expansion by tracing the development of formal accounts of analogical inference.

Genta's positive proposal takes us back to the beginning of this process. It is now only a loose fit with real cases of analogical inferences. The prediction of the material theory remains. Efforts to tighten that fit will not illuminate the value of tighter formal analysis, but will show it to be ill-advised. For it will trigger the same unproductive explosion of clauses and conditions avoided by a material analysis.

9. Jonathan Livengood
and Daniel Z. Korman, "Debunking material induction"

Jonathan Livengood and Daniel Z. Korman's contribution describes what they call the "explanatory problem of induction." It is a variant of the familiar problem of induction. According to this familiar problem, no rule of inductive inference can be justified. For all such justifications are either circular, when a rule is used to justify itself, or trigger a fanciful infinite regress of rules justified by other rules; and so on indefinitely. The variant problem proceeds from the most welcome assumption that an answer has been found through the material theory of induction to the familiar problem of induction: inductive inferences can be justified. We may have a conclusion about future occurrences well supported by justified inductive inferences. Nonetheless, the explanatory problem asserts, we lose our justification for believing the conclusion, if we come to believe that there is no suitable explanatory relationship between the conclusion and the facts it asserts. They then argue that no such explanatory relationship is provided by the material theory of induction, and, I presume, also not by any another account.

Livengood and Korman surmise my response correctly in their conclusion. There is only so much within the purview of the material theory of induction; and this problem lies outside of it. However, it is at

the border of the concerns of the material theory, so here I will delineate that border. First, the principal goal of the material theory of induction is to determine which are the relations of inductive support among propositions, where those relations are independent of human concerns and beliefs. The material analysis escapes the traditional Humeana problem of induction simply because it has no unfounded rules of inductive inference. They are all replaced by local facts or founded by local facts. However, many have suggested that an analogous regress problem resides in the circumstance that these warranting facts are in turn warranted by further facts and those in turn by still further facts. Livengood and Korman recount and accept the material theory's response, drawn from Norton (2014). There is no need for me to repeat it here.

So far, beliefs have not entered the material analysis. They enter through the historical case studies of *The Material Theory of Induction.* There, the inductive inferences of figures in the history of science are matched with the inductive inferences authorized by the material theory by, in part, reporting the beliefs of the figures. Further, the material theory is intended to be a guide to what we should belief. The theory gives no account of how we are to proceed from objective logical relations among propositions to justified beliefs. Rather it proceeded tacitly with what I take to be a near universal principle of rationality. Here is one version:

> If we are justified in believing some propositions and those propositions provide strong inductive support for a conclusion, then we are justified in believing that conclusion.

This principle and its deductive variant are tacitly supposed almost everywhere. Without it I do not see how we can sustain the standard practice of teaching elementary logic in colleges and universities. For without the principle, these logics cease to be of practical use.

The principle seems to be a truism. To be rational simply is to conform our beliefs to reason, where reason is here just a synonym for logic. Yet now Livengood and Korman bring to light deep concerns amongst epistemologists that explanatory considerations can undermine the principle. While a conclusion may conform with the logic, we may lose our justification to believe it if we come to believe that the appropriate explanatory relation to the fact concluded is absent. According to the explanatory problem of induction, there cannot be such an explanatory relation to future facts. It follows that all our rational justifications concerning future facts are fragile. Mere reflection on this impossibility is sufficient to defeat them. All I can do is urge these epistemologists to hurry up and solve the problem lest all our logical theories become useless as practical instruments for forming beliefs about future facts.

Perhaps there is a simple solution. While this is not a literature I know, none of the arguments offered in Livengood and Korman's paper seems to me sufficient to sustain this defeasibility of the principle of rationality or even to make clear precisely the explanatory relationship sought. So perhaps the principle of rationality is adequate as it stands.

10. John McCaskey, "Reviving material theories of induction"

John McCaskey's proposal that there is an extensive, overlooked, prior history to the material theory of induction is most welcome. I have a weakness for history and delight in the discovery that our latest enthusiasm is not novel. However, a closer reading of the paper shows that the overlooked tradition in inductive thought is not an earlier version of the material theory of induction, but a distinct tradition that is worthy of attention in its own right.

To see the difference, here are the salient features of the material theory of induction as I have formulated it:

(a) Inductive inference resides in relations among propositions. We infer inductively from this proposition to that; or we display relations of inductive support among propositions.

(b) The warrant for an inductive inference is provided by background facts.

(c) (My special version) All warranting facts are local.

McCaskey has identified an enduring tradition of ampliation in which the essential element is the formation of appropriate concepts. Once they are secured, ampliation – the progression from some to all – is automatic. In more modern language, the key step is identifying the projectable properties. The paper gives more details, so all that is needed here is for me to indicate how the tradition contradicts each of the features (a), (b) and (c) above of the material theory of induction.

On (a), the account does not locate ampliation at the propositional level, but at a prior level. We read (p.18, draft ms):

> I propose that we will never find such a schema for inductive inference, simply because generalization does not enter human thought at the propositional level, not at the level of sentences, judgments, and inferences. It enters at the conceptual level, at the level of words and their meanings.

On (b), what replaces my background facts (propositions) is something prior to proposition formation. It is finding good definitions. This is illustrated in the recounting of Bacon's contributions to the tradition (p. 13, draft ms):

> [Bacon] showed how, using good classification logic, well-defined concepts could lead directly to reliable, exceptionless, necessarily true, universal statements.

This is the outcome of the application of Bacon's methods (p. 16, draft ms):

> Exceptionless universal statements are then possible. Ampliation occurs at the conceptual rather than the propositional level, and classification powers generalizations.

Contrast this with my claim that background facts power inductive inference.

Finally these last quotes indicate that the locality of (c) is contradicted in so far as these procedures lead to exceptionless, universal statements or generalizations.

It will be helpful to note that McCaskey and I use the term "material" as in "material theory of induction" differently. My use is narrow, as indicated by (a), (b) and (c) above. McCaskey's use is broader. That the conceptual tradition does not employ formal schemas may well be all that is needed for him to classify it as "material."

Finally I share McCaskey's concern that we are losing sight of the importance of locating projectable properties. It is, as he notes, central to my example in Chapter 1 of Curie's inference on radium chloride. However I do understand why the approach has been marginalized. Modern problems in inductive inference are not resolved merely by the identification of projectable properties. Much more is needed to establish how the observed motions of the planets inductively support the curved spacetime of Einstein's general theory of relativity; and how the observed lines of the hydrogen spectrum inductively support the Schroedinger equation of quantum mechanics.

11. Matthew Parker, "Comparative infinite lottery logic"

Matthew Parker's paper develops an alternative inductive logic to the one described in Chapter 13, "Infinite Lottery Machines" of *The Material Theory of Induction.* The chapter assigns a monadic chance function "$Ch(.)$," to the outcomes of fair infinite lottery drawings. In its place, Parker proposes a dyadic, comparative relation among outcomes, $\preccurlyeq$, to be read as "is at most as likely as." Parker's contribution is most welcome. My goal in developing the infinite lottery logic was not merely to provide a means of assessing the strength of inductive support for various infinite lottery outcomes. Rather it was to demonstrate by

example that interesting, non-probabilistic calculi of inductive inference are appropriate when the circumstances call for it. Parker's program serves this end well. He is exploring which non-probabilistic logics are required by the facts of the infinite lottery.

His analysis draws on the existing literature in comparative probability. There has been a significant lost opportunity in that literature. Its purpose is almost never to reveal new, unexplored avenues. Rather the goal is to find that combination of properties that returns the probability calculus. It is an engaging exercise in formal mathematics, but one with a predetermined result. It is mathematically intriguing, but foundationally barren. In seeking novel alternatives, Parker's analysis exploits the untapped potential of this literature.[3]

Parker asks precisely the right question: "… what background facts about a fair infinite lottery warrant the assumption that the correct logic involves a chance function?" (Introduction). That warranting is the basis of the material theory. My analysis put most effort into justifying label independence, since that property led directly to the characteristic features of the logic. Parker is now pressing me to justify why these background facts lead to what he calls an "absolute" chance, that is a chance value assigned directly to an outcome. His alternative is that the chance relation is purely comparative.

The reasoning in support of this absolute chance is limited to judgments of equality of chance. By supposition, the defining characteristic of a fair lottery is that all that matters to the drawing of a number in some outcome set is the size (cardinality) of the set of favorable numbers in comparison with that of unfavorable numbers. If any further specification affects the drawing, then the lottery is not fair. For then some numbers are favored over others. It follows immediately that two outcomes have the same chance just if the cardinalities of their favorable sets match and the cardinalities of their unfavorable sets match.

The procedure is the same as we might use for determining equalities of chances for a fair die throw. We know that an even outcome {2, 4, 6} has the same chance as an odd outcome {1, 3, 5} since the cardinalities match. Or that an outcome {1, 2} has the same chance as {5, 6} since they too have matching cardinalities. Were the even outcome to have a different chance from an odd outcome, the die would not be fair.

Nothing more is introduced by the chance function "*Ch*(.)." To see this, consider all outcomes that arise only if one of three favorable numbers are drawn. All outcomes with just three favorable numbers have the same chance. A compact way to express this sameness is to assign these outcomes a value V_3 and declare it the value of the chance function *Ch*(.) assigned to them. This monadic chance function does nothing more than report the sameness of chance of these outcomes due to their matching cardinalities. There is no further assumption over what the value set may be.

This condition of equality is enough to give us the characteristic property of the infinite lottery logic. The chances of an even outcome ("*even*"), an odd outcome ("*odd*"), a prime number outcome, a composite number outcome and a power of ten are all equal. That is true no matter what other chance properties the lottery may have. Included in these equalities is the one Parker finds troublesome: A multiple of four outcome {4, 8, 12, 16, …}, here called "*fours*," has the same chance as an *even* outcome {2, 4, 6, …}, even though *fours* is a proper subset of *even*. This last equality is not posited independently, but it is deduced from the cardinality condition for the equality of chances.

Parker's comparative relation has the property of denying this last equality. According to it, the outcome *fours* is strictly less likely than *even*. This alternative will, no doubt, be appealing to some. For, as Parker notes (Introduction), "whenever the latter set [*fours*] wins, the former [*even*] does too, but not vice versa." The appeal derives from the obvious judgment that *even* divides into two equal sized parts: *fours* and even numbers not divisible by four. Thus *fours* is smaller than *even* and so

should be accorded a lesser chance.

The naturalness of this alternative reflects a characteristic property of finite sets not shared by infinite sets. There are two criteria of comparing the sizes of sets. Under the inclusion criterion used just now, a set is larger than any of its proper subsets. Under the cardinality criterion, a set is the same size as all those to which it can be mapped one-to-one. These two criteria agree when applied to finite sets. They no longer agree when applied to infinite sets. *Even* is strictly larger than *fours* under the inclusion criterion, but equal in size to *fours* under the cardinality criterion. This divergence is at the core of what makes the initial experience of theorizing with infinite sets disorienting.

Label independence requires judgments of equality of chances to conform with the cardinality criterion. Outcome sets of equal cardinal size, according to it, have the same chance. The results contradict the inclusion criterion and lead to the widespread violations of the condition of containment.

In devising a comparative relation that respects containment, Parker is trying to conform with both criteria. The resulting comparative relation has many appealing properties. It satisfies a comparative version of additivity, not supported by my chance function, *Ch*(.). It also respects containment, when my chance function does not. *Fours* is strictly less likely than *even*.[4] However, it is hard to serve two masters when they disagree. Efforts to do so pay a price. Parker's comparative relation secures these properties at the cost of disallowing likelihood comparisons of many outcomes sets. According to it, outcomes of even numbers and odd numbers less than 2*N* are equally likely, for all natural numbers *N*. However, the extension to infinite sets fails. We cannot compare the likelihood of even and odd outcomes without the restriction to finite subsets, no matter how natural the equality of their chances may seem. Similarly, while a *fours* outcome is strictly less likely than an *even* outcome, we cannot compare the likelihood of a *fours* outcome with an *odd* outcome, even though we might otherwise expect even and odd outcomes to be equally likely.

That Parker's relation does not compare these likelihoods is no mere oversight, but essential to the cogency of the relation. No such relation can have both:

(a) *fours* is strictly less likely than *even*; and
(b) *even* is as likely as *odd*.

While retaining label independence and transitivity.[5]

Nonetheless, it is hard to resist the appeal of a comparative relation that tells us that *fours* is strictly less likely than *even*. It would seem to be guiding us better in our lottery ticket purchases. This appearance, I contend, derives from intuitions tutored on finite lotteries. There sets larger under the inclusion criterion do have greater probability. These intuitions are a poor guide for infinite lotteries, where I do not see a formal foundation for them.

The greater probability of *even* over *fours* in the finite lottery can be expressed in terms of frequencies. Over repeated drawings, the frequencies of *even* drawings will converge towards a half and the frequencies of *fours* drawings will converge towards a quarter. That convergence is not assured. Rather it is a theorem in the probability

---

[3] For other profitable uses of this comparative notion, see Eva (2019) and Norton (2007).

[4] We have fours ⩽ even, but not even ⩽ fours, since |fours\even| = |{}| = 0, which is less than |even\fours| = |{evens not divisible by 4}| = infinity. However, we cannot compare fours with odd, since fours\odd = fours and odd \fours = odd, both of which are infinite sets. Similarly, we cannot compare even and odd since even\odd = even and odd\even = odd, both of which are infinite sets.

[5] A relabeling of outcomes renumbers fours and even as odd and not-fours respectively, so that, by (a), odd is strictly less likely than not-fours. Transitivity with (a) and (b) gives us that fours is strictly less likely than not-fours. But a different relabeling of fours and not-fours renumbers them as even and odd, respectively, so that by (b) they are equally likely.

calculus that the convergence is very probable.

There is no corresponding result for the infinite lottery. Whenever a *fours* wins, so does an *even*, but not conversely. The temptation is to shift from this result to another about frequencies: in the long run, an *even* number wins twice as often as a *fours*; or perhaps just strictly more often. As in the probabilistic case, any such convergence is not assured. It must be given to us by the chances. The chapter in the *Material Theory of Induction* demonstrates that the chances do not provide it. That is, the chance of having *n fours* drawings among *N* drawings, for each *n*, is the same as having *n even* drawings among *N* drawings, no matter how large *N*.[6] The chances give us no basis for expecting a higher frequency of *even* outcomes over *fours* outcomes in repeated drawings.

If this result is puzzling, recall that the cardinality condition operates differently for infinite and finite sets. In the finite lottery, there are twice as many numbers in the outcome *even* as in *fours*. In the infinite lottery, there are exactly as many numbers in the outcome *even* as in *fours*.

In these last remarks, I have stressed where Parker and I have differences since they might mark a way forward. In closing, I want to stress that these differences are tiny in comparison with our overall agreement on how to proceed with the chance properties of an infinite lottery. We are not to presume antecedently that chance properties must be probabilistic. Rather the chance properties are determined by the facts of the system in question. If further analysis demonstrates that these facts unequivocally determine a comparative relation such as Parker advocates, I will happily accept the result. For it would be yet another demonstration of the material theory of induction in action. The facts of a domain determine the applicable inductive logic, not our formal habits or presumptions.

12. Julian Reiss, "What are the drivers of induction? Towards a material theory"

It is comforting to find that Julian Reiss and I are in agreement on major issues concerning inductive inference. Notably he agrees with the rejection of formal theories of induction, presumably when they have aspirations of universal applicability; and allows that material facts do have a role in warranting inductive inferences. He urges, however, that the material theory of induction omits important drivers of induction, as he calls them, in restricting its warrants to facts. He lists six drivers that, he believes, have been omitted. Here I will give my reasons for disagreeing.

There is a quite general reason for doubting that there are drivers of induction independent of material facts. If one accepts that there are no universally applicable systems of inductive inference, it follows that no driver can be employed everywhere. Whether a driver can be employed in some domain is determined by the facts of that domain. That is, we trace the warrant for some inference past the driver to the facts that authorize the driver.

Reiss' six candidate drivers, however, require more individualized analysis. First are theories. The material theory allows theories to warrant inductive inferences, in so far as they are true theories. It does not align with the narrow Baconian notion of inductive inference, as Reiss suggests (Section 4.1). There are many examples of theories warranting inductive inferences in my work. Chapter 16, "A Quantum Inductive Logic," of *The Material Theory of Induction* employs the theory of quantum mechanics to warrant a particular calculus of inductive inference.

My understanding, however, is that Reiss' concern is not merely with theory, but with theories or even just postulates that have a hypothetical character. He finds an equivocation in my writing over whether facts warrant or hypotheses do. There is no equivocation here. There is one account addressing two distinct questions. What warrants an inductive inference? Answer: facts, whether the inferring agent knows the fact or

not. How can an agent know that some inductive inference is good? Answer: by learning the truth of the warranting fact. Since scientists often seek to infer while knowing too little, they hypothesize what would be the warrant for their inference if the hypothesis were to be true. Establishing its truth is necessary, but a job for future work. They thereby take on an evidential debt that must be discharged by further inductive work if they are to be sure of the security of the first inductive inference.[7]

My writing runs the two questions together. The warranting material facts were initially called "postulates" in Norton (2003) since, when we know too little, we may need to conjecture which are the applicable warranting facts. I had expected readers would be able to disentangle the two questions without trouble. Here I erred and my exposition has failed. Partick Skeels (this volume) also has had trouble distinguishing the two questions in my texts.

The next omitted drivers are idealizations and, related to them, "adequacy-for-purpose." I do allow idealizations to serve the role of a warrant. Reiss recalls one: the cosmological principle, which, construed narrowly, asserts the falsehood that the universe is exactly homogeneous and isotropic. Idealizations such as these are quite admissible as warrants, in so far as the falsities in them do no compromise the inductive inference to be warranted.

For example, the cosmological principle can be employed in inferences from the present 3 K of the cosmic background radiation to the temperatures of earlier epochs. For this inference is unaffected by the slight deviations from isotropy and homogeneity in the cosmic microwave background. If, however, the inference is to star and galaxy formation, then matters are otherwise. The slight deviations from isotropy of the order of 10^{-5} K in the cosmic microwave background reflect the inhomogenities in matter distribution that seeded star and galaxy formation through gravitational collapse.

In this context, Reiss finds our purposes to be a driver of inductive inferences. If they are to be part of the warrant for the inductive inference, then this is not so. Our purposes will tell us which inferences interest us. Whether the idealization, such as the cosmological principle, is close enough to the truth for it to serve as a warrant for the inference of interest is independent of whether our purposes make that inference interesting.

Reiss next follows a well-established literature in urging that ethical values have a role in inductive inference. His discussion recalls a familiar debate. I have nothing to add to my treatment and dissent elaborated in Chapter 5 of *The Material Theory of Induction.* Consider the proposition that there will be a planet destroying chain reaction on the explosion of the first atomic bomb near Alomogordo, New Mexico. The formulation of the proposition certainly reflects our human interests. However, once it is formulated, the strength of support provided by the evidence available to the Manhattan project scientists is independent of those human concerns. The physics of neutrons and nuclei is indifferent to the gravity of our alarm. However, that alarm will figure in how the resulting strength of support affects our actions.

Methodological norms are, for Reiss, also an additional driver of inductive inferences. His example is the choice between classical and Bayesian statistical methods. Here the material theory is unequivocal. Neither method can claim default status. They have a positive obligation in each case to justify their applicability. That justification derives from the facts of the relevant domain. This assessment has been developed at great length for Bayesian methods in roughly half the chapters of *The Material Theory of Induction.* A similar assessment would apply to classical methods. Reiss, however, asserts that facts cannot separate the methodologies in the case of the stopping rule problem, since both agree on the facts of the case. Here I differ. The facts of the case under

---

[6] Here I thank Matt Parker for assistance in the derivations of these results and for alerting me to an error in an earlier draft of them.

[7] This role of hypotheses has been developed in Norton (2014). It is not developed in The Material Theory of Induction, since that project is reserved for work presently underway.

agreement would not include those that justify the use of Bayesian or the classical statistical framework in the first place. Indeed it may even be that such facts are lacking so that one or both methodologies is unwarranted.

Finally, Reiss' sixth driver is "conceptual norms." The material theory asserts that concepts can only be used properly in an inductive inference in so far as the background facts warrant it. Take Reiss' example of the inductive inference to causes. Elsewhere (Norton, 2003a), I have argued that there is no antecedent notion of cause prior to all science. Causal talk is merely the attaching of convenient labels to processes in science, without factually restricting them. As a result, our conceptions of cause undergo continual change as they respond to factual discoveries in science. It follows that inferences to causes do not go beyond the facts of the pertinent domain. In another example, we are told that evidence favors the simpler hypothesis. There is, however, no factual, universal principle of simplicity. Rather, as I have argued in Chapters 6 and 7 of *The Material Theory of Induction*, good appeals to simplicity are really indirect appeals to particular background facts in the pertinent domain.

No doubt, the material theory of induction can be improved. Doing so by adding non-material drivers may have some initial appeal. However, since they are non-material, these additional drivers do not strengthen the theory, but dilute it and, I suspect, may even destabilize it.

13. Gerhard Schurz and Paul Thorn, "The material theory of object-induction and the universal optimality of meta-induction: two complementary accounts"

Gerhard Schurz and Paul Thorn are experts in the learning theoretic approach to induction. Elsewhere, they have provided rich and deep studies of the application of the approach. Their formal approach differs in both spirit and content from that of the material theory of induction. Hence it is encouraging to be assured by them that the two approaches agree in large measure and relate in a complementary fashion.

A substantial portion of their paper consists of a recapitulation of what they call the "optimality justification" of induction that responds to Hume's venerable problem of induction. Briefly, their goal is not to prove that some inductive method is reliable, for they accept that induction may fail. Rather, following Reichenbach, they merely seek to show that it is the best that we can do. The basis of this claim is a set of mathematical theorems on infinite sequences of possibly rounded real numbers in [0,1] and a set of functions that map initial segments to real numbers in [0,1]. The maps are "methods" or "players." They represent scientists employing inductive inference in their exploration of the world. One of them is the optimal "meta-inductivist." These collections are called "prediction games" or "possible worlds." This austere mathematical structure is intended to capture all that matters of the inductive practices of real science as far as establishing the basic theorems of the optimality justification is concerned.

Since they are the experts and have already provided a synopsis of this work in their text, there is no need for me to say more. Their synopsis should be taken as an invitation to explore their more detailed accounts and I hope that this invitation will be accepted by readers sympathetic to this style of formal analysis.

The earlier part of their paper, however, lays out in some detail where Schurz and Thorn find the material theory of induction to fail. They argue that the material theory is unable to answer Hume's problem. Their optimality justification presupposes that there is no answer to the problem better than it. Here I will explain why I disagree with their complaint.

To make their case, Schurz and Thorn seek to attribute to me what they call the "uniformity justification of induction." Warranting facts must in turn be warranted by further facts; and those by further facts; and so on. As we proceed along this sequence, they assert (Section 2, their emphasis), "the uniformity assumptions that justify material

inductive inferences become unavoidably *more and more general*." They provide an example based on my original paper of 2003. In it, they display a sequence of increasingly general warranting propositions, numbered by them (4), (5), (6) and (7). They terminate in a tight circularity. Such is the fate, they suggest, of all justificatory efforts in a material theory. (To preclude confusion, readers should be alerted that propositions (5), (6) and (7) are not drawn from my work but are conjectures over how a material theorist would proceed. I do not endorse them.)

This material theorist sees the problem quite differently. In order to produce the troublesome regress, Schurz and Thorn make an assumption about the large-scale structure of relations of inductive support. They attribute a hierarchical structure to it in which propositions of one generality require warrants in propositions of greater generality. An examination of actual relations of support in science fails to shows this hierarchical structure. Instead we find a massive tangle of relations of support admitting no such hierarchy. No regress of the type envisaged by Schurz and Thorn is found there.

Schurz and Thorn seek the justification for induction within the material theory at the end of regress they construct; and of course they do not find it. In the material theory of induction, there is no single locus, like the end of a rainbow, where the justification of inductive inference as a whole is found. The justification is distributed over the entirety of the relations of inductive support. Pick any proposition central to a mature science. The justification for it can be displayed. There is no such proposition for which a justification cannot be given. One can, if one wants, try to trace out the relations of support. Because of the non-hierarchical and tangled nature of the relations of support, one rapidly takes a tour through and round an enormous maze, whose paths divide and divide again, and perhaps some even lead back to the starting point. None end poorly.

This material analysis provides all the justification needed and in a way that is recognizable in actual science. There is no basis for doubt of the periodic table of the elements or any other of the myriad of propositions in a mature science simply because each is well-supported. That support resides in the relations among the various propositions and does not require any of the general formal schemas whose adoption leads immediately to a fanciful regress or the dubious, tight circularities of the traditional Hume problem. To accept that fact, but still to harbor doubts is to ask for justification while ignoring its presence. It is akin to a novice engineer who examines the individual beams and struts of some fanciful architectural structure and find each to be well-supported. Nonetheless the engineer somehow still doubts that the structure can stand.

This analysis has been developed in greater detail in Norton (2014), which unfortunately does not appear in Schurz and Thorn's list of references. That it can be overlooked is understandable. I avoided all talk of it in the chapters of *The Material Theory of Induction* until that avoidance was finally explained briefly in the work's epilog, where I foreshow further elaborations now in preparation.

14. Patrick Skeels, "A tale of two Nortons"

Patrick Skeels' "A Tale of Two Nortons" will be alarming to those who find one Norton to be already one too many. The first Norton encountered by Skeels argues that inductive inferences are justified by facts. The second Norton argues that inductive inferences are justified by knowledge of those facts. It is, we are told (Section 1.3.2), "less than clear what position Norton is, in fact, defending"; and (Introduction) "one may worry that Norton is vacillating between two substantially different views."

Fortunately there is only one Norton writing on the material theory of induction and only one theory. There are however two questions. Skeels has mistaken the treatment of two questions by one theory as the implementation of two theories. The two questions are:

(inductive-logical)

Question: Which inductive inferences are good?

Answer: Those that are warranted by a (true) fact.

(epistemic)

Question: How can we know that a specific inductive inference is good?

Answer: We must be assured of the truth of the appropriate warranting fact.

The first question pertains to logical relations among propositions. The warranting facts do their work independently of human knowledge and awareness. Nothing precludes awareness of them. It is just that their warranting powers are independent of such awareness. The second pertains to the processes of gaining this knowledge. The answers to both questions are closely coupled. The warranting fact of the answer to the second question is the fact that warrants in the answer to the first question. The two Nortons of Skeels' critique are really just one Norton answering these two questions with one theory.

In answering them, I did not make any real effort to distinguish the questions sharply. I thought the context quite sufficient to discriminate them. It was of some concern to me on reading Skeels remarks that at least one reader was confused. I reread the two papers on which Skeels concentrated, Norton (2003, 2014), and I was reassured that a sympathetic reader would have little trouble recognizing the distinctness of the questions and the unity of the answers.

Skeels has quoted passages symptomatic of the two Nortons. The "fact" version of Norton is quoted (Section 1.1) as saying.

> [Facts] 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.

This passage answers the "inductive-logical" question. It makes a point in inductive logic, distinct from any matters of human knowledge and awareness. The "knowledge" version of Norton is quoted by Skeels (Section 1.3.1) as saying:

> 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].

This second remark clearly addresses the discovery activities of inferring agents who seek to learn. It answers the epistemic question. It does so by drawing on the answer to the inductive-logical question. This may be less clear to readers of Skeels since his quote omits the words "the material postulate that licenses the induction" in the original passage, included here in brackets.

There are many more pertinent remarks, supposedly by the two Nortons.[8] I am confident that, with this clarification, readers will have little trouble seeing that they are remarks by one Norton with one theory on two problems.

While it is tedious to dispute minor points, one such remark requires a correction. Skeels remarks in his introduction that "Norton's material theory was developed for the explicit purpose of dissolving Hume's Problem …" As a simple matter of biography, this is not so. The ideas of the original 2003 paper were fully worked out before Jim Bogen[9]

pointed out that there was a connection to the problem of induction. The problem occupies only 3 pages of the paper. My monograph, *The Material Theory of Induction*, contains 16 chapters, none of which mention the problem of induction. It is mentioned only in an Epilog, when I explain why it was not mentioned earlier. The reason is that the problem provides a distraction so irresistible to many philosophers that they then fail to see that the real goal of the material theory of induction is to show the material nature of inductive inference. That the theory dissolves the problem of induction was an unexpected bonus and a welcome one at that.

The reason this minor misunderstanding requires correction is that it forms the basis of Skeels' critique of the material theory of induction. His stated purpose (Section 3) is to show that "both [Nortons] failed to accomplish their intended purpose of dissolving Hume's Problem." And, as a result, "If my arguments are successful, then the material theory faces a significant difficulty …. It is unclear whether or not the material theory is beyond repair …" I do not believe that the material theory of induction is in need of any repair. A failure to dissolve the problem of induction does not impinge on the theory's goal. It would merely put it in the company of very many other such failures.

That said, it does not seem to me that Skeels' critique impugns the dissolution. I stand by my account and invite readers to read it and make up their own minds. It will be useful, however, if I indicate here why I am unmoved by Skeels' critique.

While his reasons likely differ from mine, Skeels accepts (Section 2.1.1) that the problem of induction is dissolved in the context of the inductive-logical question above:

> It must be admitted that the fact version of the material theory does appear to solve Hume's Problem. Indeed, it neutralizes it altogether.

However, he continues to argue that a "formal theorist" can proceed by "claiming that facts can justify formal inductive inferences in precisely the same way." From this he concludes that the material theory is "completely immaterial" to the dissolution and the view is "somewhat self-defeating."

This is a gratifying concession since it grants that the material theory has actually dissolved the problem of induction in the inductive-logical sense. I also welcome formal theorists who decide to justify their schemas with material facts. In so doing, they have become material theorists. That is not a defeat, but victory for the material theory.

I have argued that the relations of inductive support in mature sciences form a massively tangled structure that is self-supporting. While circularities are inevitable and rampant in this structure, I believe the circularity is benign. In this aspect, my account is similar to coherentist epistemology, although I am not a coherentist epistemologist. Skeels is willing to grant that the circularity in coherentist epistemology is benign. However he refuses to grant it to the material theory of induction. His reason (Section 2.2) is that beliefs can enter into benign circularities, he says, "precisely because they have content, e.g. propositions can entail other propositions but they cannot entail a rule or an action." He couples this with the claim that material theory concerns rules for inductive inference. Hence "coherence simply is not applicable when they [rules] are present."

What Skeels overlooks with this objection is that propositions can also be equivalent to rules in the sense that their meaning authorizes inferences. In deductive logic, "If A then B″ is both a proposition and authorizes a rule that allows one to infer from A to B. In inductive logic, the proposition "samples of elements are generally uniform in their properties" is both a fact and also authorizes inductive inferences among the properties of samples of elements. Once this dual function is recognized, Skeels' objection fails.

Finally, "Norton's Decision" is that agents are supposed to be in a quandary over whether to make an inductive inference, since justification is blocked by Hume-problem-like circularities. It is presented as a new and unanswered question. That is quite puzzling since the challenge

---

[8] I found only one exception in the two papers, Norton (2003, 2014), in which I simply misspoke ("miswrote"). Norton (2003, p. 658) says "We can see immediately that the material postulate that underwrites our inference in accepting these result is just our belief that the method is reliable." I collapsed two points into one. I meant, first, that the inference is warranted by the fact of the reliability of method; and, second, that our acceptance of the inference derives from our belief in this fact.

[9] He is acknowledged for the point in Norton (2003, p. 667, footnote 8).

is just the "epistemic" question above, when set within the context of my treatment of the problem of induction. The answer I have given is that agents can know that an inductive inference is justified by tracing out its warrant in the complicated, non-hierarchical tangle of relations of inductive support described in my papers.

Skeels overlooks this answer. The reason seems to reside in his remark:

> … we lack direct, unmediated, epistemic access to the facts. Once we separate justification from decision in the way that the fact version of the material theory does, this question becomes especially pertinent as it cannot be solved in the same way as Hume's Problem.

The difficulty here is obvious. It is only Skeels' misreading of the material theory that precludes an agent having access to these facts.

It seems that this misreading arose through Skeels identification of the "fact" Norton as an external epistemologist. It is against this externalist "fact" Norton that the decision problem is directed.[10] In an externalist epistemology, agents need not and may not have access to the justifications of their inferences. Thus the "fact" Norton can suppose that agents have no access to the facts that warrant their inductive inferences. This preclusion is simply an artifact of Skeels' imposition of an externalist epistemology onto this part of the material theory. There is nothing in the theory that precludes an agent identifying the warranting facts of an inductive inference if the agent seeks it. I often write of agents doing just this.

Indeed, I am left wondering if the creation of the two Nortons is a result of misguided efforts to divide the claims of the material theory into mutually incompatible externalist and internalist sets. Those efforts create the externalist "fact" Norton and the internalist "knowledge" Norton. Since externalist and internalist epistemologies, so construed, are incompatible, these invented Nortons must harbor incompatible commitments. The resulting misreading of the material theory would then be a lesson in the dangers of imposing ill-matched categories onto a theory.

15. Michael Stuart, "The material theory of induction and the epistemology of thought experiments"

Michael Stuart's contribution to this volume gives us more than a critical response to the material theory of induction. It is also a proposal for an extended epistemology of thought experiments that can accommodate the material conception of inductive inference. It offers new and interesting insights into thought experiments and is well worth careful study.

The motivation Stuart provides for his account is a tension in my work in two areas. In my analysis of the epistemology of thought experiments, I argue that thought experiments are simply picturesque arguments. This analysis was completed in its major conceptions before I worked on the material theory of induction. In the thought experiment analysis, I proceeded with a formal characterization of inference, both deductive and inductive. The material theory now disputes this characterization of inductive inference. Stuart is correct to identify the tension and urge the need for a reconciliation. It is this aspect of his paper that I will discuss here.

Stuart is greatly troubled by the tension and finds it to require some major concessions from my original argument account of thought experiments. Those concessions conveniently open a space for the embellishments of his own account. My assessment, however, is that substituting a material conception of inductive inference requires rather little to change in the argument account. Indeed, in some areas, it is strengthened. Stuart has described seven theses in the argument

account. They are.

> Identity Thesis, Reconstruction Thesis, Reliability Thesis, Elimination Thesis, Empirical Psychology Thesis, Empiricist Thesis.

All seven are retained, in my view, after the material reconception of inductive inference.

In so far as the argumentation in a thought experiment is deductive, then no modification is needed. These deductive thought experiments constitute a significant portion of the whole corpus of thought experiments in science and contain many of the most prominent examples from all eras. They include Stevin's sixteenth century looped chain thought experiment; Maxwell's original nineteenth century demon thought experiment; and Einstein's thought experiments that establish the relativity of simultaneity and $E = mc$.[2]

Changes are needed only when the argument in the thought experiment is inductive. Then there is one casualty. In Norton (2004, pp. 52–53) I argue that the reliability of a thought experiment depends on our being able to identify a "mark" within it that is then identified as some formal feature associated with a general notion of a logic. To accommodate material inductive inference, this argument must be generalized to include marks derived from warranting background facts. I expect this can be done, since background facts can warrant mini-logics applicable just to their specific domains. Will the resulting modified argument be successful? Answering awaits someone willing to undertake a full working out of the needed modification.

Otherwise, the claims and arguments of the account stand. I revisit a few of them here. The principal thesis ("Identity Thesis") is that thought experiments are just picturesque arguments. My main argument for this derives from the premise that "pure thought cannot conjure up new knowledge" (Norton, 2004, p. 50). All that the armchair reflections of thought experiments can do properly is to transform what we already know in a way that preserves its truth (deductive inference) or preserves its likelihood (inductive inference) (p. 49).

A second and lesser[11] argument (p. 50) is that there are no examples known to me of well-functioning thought experiments in science that cannot be reconstructed as arguments. Stuart is correct to identify this inference from cases to the general claim as an inductive inference and he is correct to ask me for the background facts that warrant it. Stuart (Section 4) conjectures without success what these background facts might assert of thought experiments. He is searching in the wrong place. The inductive inference is at the metalevel of the thought experiment literature and that is where these background facts are found. They are that the total corpus of scientific thought experiments is small and well-explored by philosophers. Over the last 30 years, the argument account has been a prominent target of criticism. Yet, in spite of the extended efforts of critics to find counterexamples, none has emerged. Indeed, if thought experiments could open a novel epistemic channel that transcends the reach of argumentation, there would be a powerful incentive for scientists to employ it. Yet no such thought experiment has been located by the philosophical literature.

The cases made for other theses in the argument account remain. An example is the Empirical Psychological Thesis, which Stuart wants to discard. The thesis asserts that the actual conduct of a thought experiment is the execution in thought of the reconstructed argument. This follows from a straightforward reading of common thought experiment texts: the narrative simply walks the reader through the steps of an argument. It also follows from the fact that the execution of a thought experiments gives us nothing more than the conclusions of the reconstructed argument. If some mode other than argumentation is at work, that mode has the curious property of mimicking argumentation perfectly in what it can tell us. This justification remains when we

---

[10] "Unlike externalists, internalists do not have to face Norton's Decision …" (Section 2.1.2.).

[11] Unfortunately, Stuart (Section 4) misidentifies this second argument in two forms B) and C) as the "main pieces of evidence for the argument view."

replace a formal conception of inductive inference by a material conception.

Finally, the argument given in Norton (2004, pp. 52–53) for the reliability thesis requires modification. However, the reliability thesis persists as a challenge to any candidate account of thought experiments. As Norton (2004, §2.2) shows, there are many cases of thought experiments whose results are contradicted by a second thought experiment. What independent provision does the candidate account provide to enable us to separate the good thought experiments from those that mislead? If thought experimenting is to be a reliable instrument, there must be some way of doing this. No other account of thought experiments in the survey of Norton (2004) can do this except through the use of some argument structure. The argument account effects the separation by assessing the cogency of the argumentation.

Replacing a formal notion of inductive inference with a material notion strengthens the argument account since the material approach gives much better control of inductive inference. To use Stuart's example, Newton and a composite Mach-Einstein arrive at competing conclusions in the Newton's bucket thought experiment. Each is inferred as a best explanation. If we approach inference to the best explanation formally, our adjudication will require some general assessment of who truly explains better. The competing and convoluted accounts of explanation in the present literature can provide no clear decision at this general level. However, the material analysis of inference to the best explanation in Chapters 8 and 9 of *The Material Theory of Induction* will guide us quite quickly to the differences of background facts presumed that lead to their different conclusions.

A second example concerns the use of typicality as a form of inductive generalization in thought experiments. Assessing the strength of such typicality inferences by general, formal means leads to intractable vagueness. A material analysis treats each such inference as, in principle, distinct and in need of its own warranting facts. Identifying them, or their absence, enables appraisal of the cogency of the argument.[12]

My account of thought experiments seeks only to solve what I call the "epistemological problem of thought experiments in science" (Norton, 2004, p. 44):

Thought experiments are supposed to give us knowledge of the natural world. From where does this knowledge come?

Stuart's "pluralist epistemology" goes beyond this problem and considers, among other things, how we might create[13] a thought experiment, its broader aims, the epistemic virtues of reasoning agents and at least mentions their rhetorical prowess. These bring welcome and fruitful expansions of the original epistemological problem, but they require no further corrections to my original argument account of thought experiments.

References

Eva, B. (2019). Principles of indifference. *Journal of Philosophy, 116*, 390–411.

Norton, J. D. (2003). A material theory of induction. *Philosophy of Science, 70*, 647–670.

Norton, J. D. (2003a). Causation as folk science. *Philosophers' Imprint, 3*(4). (Reprinted in pp. 11–44, H. Price and R. Corry, Causation, Physics and the Constitution of Reality. Oxford: Oxford University Press).

Norton, J. D. (2004). Why thought experiments do not transcend empiricism,. In C. Hitchcock (Ed.), *Contemporary debates in the philosophy of science* (pp. 44–66). Blackwell.

Norton, J. D. (2007). Probability disassembled. *The British Journal for the Philosophy of Science, 58*, 141–171.

Norton, J. D. (2008). Cosmic confusions: Not supporting versus supporting not-. *Philosophy of Science, 77*, 501–523, 2010.

Norton, J. D. (2014). A material dissolution of the problem of induction. *Synthese, 191*, 671–690.

Norton, J. D. (2018). The Worst thought experiment. In M. T. Stuart, J. R. Brown, & Y. Fehige (Eds.), *The routledge Companion to thought experiments* (pp. 454–468). London: Routledge.

Stoney, G. J. (1871). On the cause of the interrupted spectra of gases. *Philosophical Magazine, 41*, 291–296.

Tutton, A. E. (1910). *Crystalline Structure and chemical composition.* London: MacMilllan & Co.

---

[12] For an example of typicality inferences in a thought experiment that has caused long-standing mischief, see my account (Norton, 2018) of the "The Worst Thought Experiment."

[13] To preclude confusion, Stuart (Section 7.1) associates this discovery of the thought experiment with the "context of discovery." I also use the term (Norton, 2004, p. 50) to refer to activity of carrying out the argument in an existing thought experiment to discover it conclusion.