

Debunking Material Induction

Jonathan Livengood

University of Illinois at Urbana-Champaign

Daniel Z. Korman

University of California Santa Barbara

Word Count (including references and notes): 8087

Abstract. In this paper, we present an explanatory objection to Norton's material theory of induction, as applied to predictive inferences. According to the objection we present, there is an explanatory disconnect between our beliefs about the future and the future facts. We argue that if we recognize such a disconnect, we are no longer rationally entitled to our future beliefs.

Keywords: Induction, Norton, Material Theory, Debunking, Explanatory Constraint

Acknowledgements: We would like to thank Kathleen Creel, Sam Fletcher, Kasey Genin, Dan Malinsky, Conor Mayo-Wilson, Jonah Schupbach, Kieran Setiya, Elay Shech, Tom Sterkenberg, and Kino Zhao for helpful and challenging comments on earlier drafts. We would also like to thank John Norton for several email exchanges about his views during the preparation of our paper and for his own extensive and provocative writings on induction, which we greatly admire.

Debunking Material Induction

Induction is indispensable in the sciences and in daily life. Yet, it continues to be a scandal in philosophy. Over the last 15 years, John Norton has been developing and defending an account of induction that substantially reduces its mysteriousness and makes its justifying power easier to see. His key insight is to treat the goodness of inductive inferences as a *local* matter. On Norton's *material theory of induction* (Norton 2003, 2010, 2011, 2014, and forthcoming), there is no globally-applicable, formal rule of induction, and no universal fact that underwrites all good inductive inferences. Rather, a good inductive inference is powered by local matters of fact specific to that inference or to a restricted family of similar inferences.

Of course, Norton is not without critics,¹ but we think that he has adequately addressed the standard objections to his theory. Moreover, we think that the material theory throws substantial light on our inductive practices and is enormously useful for inquiries in the history and philosophy of science.² Nevertheless, we think that lying at the heart of Hume's skeptical attack on induction is a challenge that the material theory is not equipped to handle. In this paper, we advance an explanatory problem of induction and argue that our explanatory problem threatens even accounts (like Norton's) on which there is no universal, formal rule of induction.

Here is how we proceed. In Section 1, we sketch a standard formulation of the problem of induction and show how the material theory is meant to solve it, and in Section 2 we review Norton's responses to a common objection to the solution. We then lay out our explanatory formulation of the problem of induction in Section 3, and we address some candidate responses to the problem in Sections 4-5.

¹ For examples, see Achinstein (2010), Kelly (2010), Steel (2005), and Worrall (2010).

² Cf. Livengood and Edwards (forthcoming).

1. The Material Theory of Induction

To illustrate the difference between formal and material theories of induction, consider the following two inferences:

[T1] On all past occasions, there has been exactly one low tide on the Thames at London Bridge during daylight hours.

T =====

[T2] Tomorrow, there will be exactly one low tide on the Thames at London Bridge during daylight hours.

[B1] On all past occasions, samples of bismuth have melted when heated to 271.4° C.

B =====

[B2] On the next occasion, bismuth will melt when heated to 271.4° C.

Both [T] and [B] are valid inductive inferences.³ On a formal theory of induction, what makes these inferences valid is that they conform to a rule of inductive inference. Perhaps both inferences are instances of the same universal formal rule (such as a rule that allows one to infer that some property will be instantiated on the next occasion from the fact that it has been instantiated on all past occasions) or perhaps there is a different rule that underwrites each inference. On the material theory, what makes [T] a valid inference is the fact that the Thames is an estuary river and estuary rivers generally have regular tides. What makes [B] a valid inference is that bismuth is an element and a sample of a given element generally has the same melting point as other samples of that element. In both cases, local material facts power the inferences.

In addition to the question of what makes an inductive inference valid, there is a question of when one is *justified* in drawing the inference. Norton isn't explicit about this, but we think

³ We follow Norton (2014, 672) in using “valid” to characterize inductive inferences as well as deductive inferences.

the right thing for the material theorist to say is that one is justified in drawing such inferences only if one is justified in believing that the underwriting material facts obtain.⁴ After all, if one has no reason to believe that samples of a single element generally have the same melting point, one wouldn't be justified in inferring [B2] from [B1].

To see the material theory in action, let's consider an argument for inductive skepticism that it's tailor-made to handle. Let a *predictive* inference be an inductive inference about how things will be in the future.⁵ Let the Principle of Induction [PI] be the proposition that the future will be like the past.⁶ Then the skeptic argues as follows that no predictive inference can be justified:

- [S1] [PI] cannot be justified.
- [S2] If so, then no inference that presupposes [PI] is justified.
- [S3] Every predictive inference presupposes [PI].
-
- [S4] No predictive inference is justified.

The usual motivation for [S1] is that, since [PI] is neither a priori nor directly given in experience, it would have to be justified on the basis of an inductive inference. But any such inference would either have to be underwritten by [PI] itself, rendering the inference viciously circular or have to be underwritten by some further principle that would itself have to be underwritten by some still further principle (on pain of circularity), and so on *vitiose ad*

⁴ Cf. Kelly (2010, 759-760).

⁵ See Salmon (1981), *contra* Popper (1972), on the centrality of such inferences to scientific inquiry.

⁶ Alternatively, one might characterize the Principle of Induction in terms of the idea that nature is uniform. We don't think anything that we say hangs on going one way rather than the other.

infinitum. Premiss [S2] is a special case of the more general principle that no inference with an unjustifiable presupposition is justified.

As for [S3], one can see why proponents of formal theories of induction seem compelled to accept it (or something like it). By their lights, [PI] (or some closely associated rule of inference) is what licenses the move from propositions about the past and present to propositions about the future in predictive inferences like [T] and [B] above. According to the material theory, by contrast, [S3] is false. [T] is underwritten by a local material fact about tides, [B] is underwritten by a local material fact about elements, and neither presupposes anything as general as [PI]. Indeed, the proponent of the material theory can happily accept that [PI] is far too vague (as stated) to explain which inductive inferences are and aren't valid, and is spared the embarrassment of having to identify some more precise or nuanced general rule of induction that won't (for instance) license the conclusion that there will never be a cure for cancer or that we will never turn 60 (since on all our past birthdays...).

2. Resisting a Material Regress

Some may naturally wonder whether the material theory faces a regress of its own. As we saw, an inference like [B] is justified on the material theory only if one is justified in believing the material fact that underwrites it: namely, that samples of a given element generally have the same melting point as other samples of that element. Following Norton, let's call such claims about how things generally are "material postulates," and let's call this particular material postulate [B*]. What justifies us in believing [B*]? It would have to be some further material postulates, for instance that samples of a given element generally have the same microscopic structure and that things with the same microscopic structure generally agree in their physical properties (see Norton 2014, 676).

Now the regress is in motion. For one can ask what justifies us in believing those further material postulates and what justifies the postulates that justify them and so on. The chain of material postulates can't go on forever. So there would seem to be two possibilities. The first is that it circles back on itself, with some postulate underwriting the very inferences that ultimately justify us in accepting that very postulate. The second is that the chain of postulates terminates in some collection of first postulates, with no prior postulates to justify them. But if the first postulates aren't themselves justified by postulates, then they aren't justified. After all, the postulates that power inductive inferences are neither a priori nor directly given in experience, and therefore, they can only be justified by inductive inferences which (by hypothesis) must be powered by postulates.

Norton addresses such worries, rejecting the underlying “hierarchical” conception of justification.⁷ But rather than appealing to a Quinean web to illustrate the holistic nature of justification, his preferred metaphor is that of a stone arch, with each stone figuring in the explanation of what holds the others in place (2014, §10). Crossword puzzles can serve the same illustrative purpose. When a completed puzzle yields sensible answers to all of the clues, each answer is justified at least in large part by its coherence with the rest of the answers—even if the answer to any one clue, taken in isolation, is only a guess.⁸

The material postulates that figure in the justification of inductive inferences likewise, by Norton's lights, form “an inductively self-supporting system of propositions” (2014, 686). He gives numerous examples of nonhierarchical “cross-overs”:

When we infer inductively from the observed positions of Venus to the elliptical orbit that fits them, we select an ellipse from all possible curves, on the warrant of Newton's inverse square law of gravity that entails that planets orbit in ellipses. Yet Newton's law

⁷ He also challenges the claim that postulates cannot be directly justified by experience, as resting on the myth of the given; see his 2014, §9.

⁸ Haack (1993, Chapter 4) develops the crossword puzzle analogy in detail. See also, Huemer (2001, 40-41).

can have the benefit of inductive support from the very ellipses whose fitting it warrants. (2014, 684)

Multiple crossings over of support relations are a commonplace of science. It is quite visible in the interactions among different sciences, both in general and narrow propositions. Each science, for example, is confident that the general principle of the conservation of energy holds within its domain. It draws that confidence, in part, from knowing that energy conservation holds within the other sciences. (2014, 686)

In the case at hand, then, the idea would be that what justifies us in believing [B*] is its coherence with a wide range of further postulates; and the justification for any one of those other postulates is its coherence with the rest, including [B*] itself.⁹

We think Norton's response to the regress objection is satisfying, though we might quibble with some of the details. We want to be clear about this because one might be tempted to read our explanatory objection as a variation on the regress objection. To preview what we say in the next two sections, the explanatory objection is importantly different from regress objections both in terms of what it concedes to Norton and in terms of the challenge it raises. The explanatory objection concedes that when people draw predictive inferences, they are typically justified in exactly the way the material theory suggests. We are not demanding a *justification* in a way that leads to a regress. Rather, the explanatory objection advances a defeater that, once brought to our attention, undermines whatever justification we might otherwise have for our beliefs.

3. An Explanatory Constraint on Rational Belief

In this section, we lay the groundwork for an explanatory problem of induction. Our formulation of the problem is inspired by our reading of Hume, which is similar to Jacobson's

⁹ One may naturally wonder whether the appeal to coherence would be equally effective in vindicating formal theories. Norton will likely deny that it would: his "arch" is a completed structure (viz. mature science) whose parts clearly support each other, whereas the formal theorist has little more than three rocks arranged hopefully (viz. a small supply of ill-defined formal rules). Cf. his 2014, 676.

(1987) reading, and by recent debunking arguments against moral realism, for which see, e.g., Street (2006) and Joyce (2006).¹⁰ The problem is generated by an explanatory constraint on rational belief, which we formulate as follows:

[EXC] If *S* believes that there is no appropriate explanatory connection between her belief that *p* and the fact that *p*, then *S* is not rationally entitled to believe that *p*.

In other words, recognizing that some beliefs lack an appropriate explanatory connection to the facts that they purport to be about serves as a defeater for those beliefs. (We will say more later on about what might count as an “appropriate” connection.)

Some may reject the need for any such constraint out of hand. Indeed, they may reject it precisely on the grounds that it seems to preclude predictive inferences.¹¹ There are, however, several excellent reasons for thinking that there must be some such explanatory constraint on rational belief, which we’ll survey momentarily. But before motivating [EXC], let us get clearer on what it does and doesn’t demand.

First, [EXC] does not demand that anyone positively *believe* that there is an explanatory connection between their beliefs and the relevant facts, thus avoiding the implication that very young children—who do not yet have any beliefs about explanation—are not rationally entitled to any of their beliefs. Second, [EXC] does not have the implication that one has to establish that there is an explanatory connection or spell out how such connection works in order to believe rationally. Our ancient ancestors, lacking any good explanation of how perception works, were still rationally entitled to their perceptual beliefs, since they (presumably) didn’t positively deny that worldly facts explained their perceptual beliefs. Third, [EXC] does not demand that *there actually be* an explanatory connection between one’s beliefs and the facts. Accordingly, it does

¹⁰ See Setiya (2012) and Dogramaci (2016) on the connection between moral debunking and inductive skepticism.

¹¹ See, e.g., White (2010, 582-583) and Fairchild and Hawthorne (2018, 52).

not entail that brains in vats lack rational entitlement to their external-world beliefs. What [EXC] does say is that for an epistemic agent to have rational entitlement to believe a given proposition, she cannot believe that her belief is disconnected from the relevant fact.¹²

Why accept [EXC]? We find the idea of an explanatory constraint on rational belief obvious and compelling. But there are several ways to motivate an explanatory constraint that should appeal even to those who do not find it so obvious and compelling. In the rest of this section, we canvas three reasons to accept—or at least take very seriously—the claim that there is an explanatory constraint on rational belief substantially similar to [EXC]: first, an explanatory constraint on rational belief may be motivated by way of an explanatory constraint on evidence; second, an explanatory constraint on rational belief may be motivated by way of non-accidental-connection constraints on knowledge; and third, an explanatory constraint on rational belief may be motivated by way of accounting for the epistemic force of access challenges.¹³ We expect that each way of motivating an explanatory constraint on rational belief could be developed in several different ways. For each of the three routes, we try to provide a clear, concrete way of arguing for [EXC], but we hope that if our specific arguments are found wanting, the general ideas will remain plausible.

Let's begin by considering an explanatory constraint on *evidence*. A version of such a constraint may be found in Section 4 of Hume's *Enquiry*. After giving some examples of inferences to "absent facts," Hume writes (Section IV, Part 1, Paragraph 22, Pages 26-27):

All our reasonings concerning fact are of the same nature. And here it is constantly supposed that there is a connexion between the present fact and that which is inferred from

¹² Although "disconnected" is not strictly-speaking synonymous with "not appropriately connected," if a belief is not appropriately connected to the relevant fact, we will sometimes say that the two are disconnected.

¹³ In addition to the three reasons we discuss, we think that an explanatory constraint on rational belief provides a compelling diagnosis of what is wrong with some intuitively objectionable theoretical maneuvers, such as level-splitting with respect to higher-order evidence (for discussion of which, see Horowitz, 2014). But we do not develop this further reason here.

it. Were there nothing to bind them together, the inference would be entirely precarious. The hearing of an articulate voice and rational discourse in the dark assures us of the presence of some person: Why? because these are the effects of the human make and fabric, and closely connected with it. If we anatomize all the other reasonings of this nature, we shall find that they are founded on the relation of cause and effect, and that this relation is either near or remote, direct or collateral. Heat and light are collateral effects of fire, and the one effect may justly be inferred from the other.

The idea seems to be that for any given inference, if our evidence had no appropriate explanatory connection to what we infer from it, then the inference would only yield a true conclusion by chance. The success of an inference that drew a true conclusion would be unstable, fragile. Hence, disconnected evidence provides no reason for belief.

Achinstein (2008, 344-345) develops and defends an explicit explanatory constraint on evidence beginning with the idea that the function of evidence (especially in the sciences) is to provide us (scientists) with good reason to believe a hypothesis. Put more precisely, Achinstein's *good-reason-to-believe assumption* says that e is evidence that h (given b) only if (given b) e provides a good reason to believe h . After discussing some attempts to satisfy the good-reason-to-believe assumption and showing how they fail, Achinstein proposes a principle that satisfies the good-reason-to-believe assumption by tying together (1) the fact that e is a good reason to believe h and (2) the probability of there being an explanatory connection between h and e .¹⁴ The result is an explanatory constraint on the nature of evidence: some proposition e is evidence for some other proposition h only if there is a high enough probability (greater than $\frac{1}{2}$) that there is an explanatory connection between h and e , given e and the rest of our background beliefs (346).

Our explanatory constraint on rational belief does not immediately follow from Achinstein's explanatory constraint on evidence. But the two are related by way of an *epistemic*

¹⁴ Regarding explanatory connection, Achinstein writes (346), "There is an explanatory connection between h and e if and only if h correctly explains why e is true, or if e correctly explains why h is true, or if some hypothesis correctly explains both why e is true and why h is true."

instrument principle, according to which if *S* has good reason (all things considered) to believe some proposition *p*, then the fact that [*S* believes that *p*] is evidence that *p* is the case.¹⁵ Now, if the fact that [*S* believes that *p*] is explanatorily disconnected from the fact that *p*, then according to Achinstein's constraint, the fact that [*S* believes that *p*] is not evidence that *p*. If the fact that [*S* believes that *p*] is not evidence that *p*, then it follows from the epistemic instrument principle that *S* does not have good reason (all things considered) to believe that *p*. And if *S* comes to believe that she does not have good reason (all things considered) to believe that *p*, then she has a defeater for her belief that *p*, just as [EXC] maintains.

A second reason to take an explanatory constraint on rational belief seriously is that an explanatory constraint on rational belief follows from a non-accidental-connection constraint on knowledge.¹⁶ Knowing that *p* requires that it not be an accident that one correctly believes that *p*.¹⁷ Many philosophers have urged that non-accidental connection should be understood in modal terms.¹⁸ But Setiya (2012, 89-99) and Faraci (2019, 10-15) argue that the operative sort of accidental connection cannot be cashed out that way and must instead be understood in explanatory terms. If that's right, then a no-accidental-connection constraint on knowing amounts to an explanatory-connection constraint on knowing. To get from there to [EXC], two things are needed.

¹⁵ The epistemic instrument principle could be stated in terms of the reliability of the epistemic agent, but we worry that appealing to reliability in this context will naturally suggest a modal account that we do not endorse. Probably, the principle should also be restricted to agents other than *S*, but that restriction doesn't matter to our argument here.

¹⁶ One way of spelling out a non-accidental-connection constraint is in terms of causal connection. One of the authors of this paper thinks that causal accounts of knowledge are still worth considering, despite the fact that the literature seems to have moved on. If you think that there is something right about a causal account of knowledge, then notice that on a causal account, if you come to recognize that a belief is causally disconnected from the fact believed, then you gain a defeater for that belief, just as you would if you recognized that some other condition on knowledge (such as truth) was not met. But since appropriate causal connection is a kind of explanatory connection, [EXC] follows immediately from a causal account of knowledge.

¹⁷ Unger (1967) claims that being non-accidentally connected to the fact believed is a necessary condition for a belief to count as knowledge. Unger (1968) claims that knowledge just is non-accidental belief.

¹⁸ Pritchard (2005, 2007, 2009) is probably the most prominent example. See Faraci (2019) for more examples.

First, we need the following plausible principle connecting constraints on knowledge and constraints on rational belief:

For all constraints K on knowing and all propositions p , if you come to believe that your belief that p does not satisfy K , then you have a defeater for your belief that p .

Second, we need a more precise formulation of explanatory construal of accidental connection: one's belief that p is accidentally true if there is no explanatory connection between the belief that p and the fact that p .¹⁹ Together with a no-accidental-connection constraint on knowledge, these yield the result that if you come to believe that one of your beliefs is not explanatorily connected to the fact believed, then you have a defeater for that belief. But if you have a defeater for the belief that p , then you are not rationally entitled to believe that p , just as [EXC] says.

A third reason to accept that there is an explanatory constraint on rational belief is that such constraints play an indispensable role in well-known and *prima facie* forceful access challenges in the philosophy of mathematics (due to Benacerraf 1973). We know that three is odd and greater than two. But if—as Platonists would have it—this is a fact about the features of, and relations among, intangible, mind-independent, and causally inert objects, it's difficult to see *how* we know that three is odd and greater than two. The relevant facts are inaccessible on a Platonist view. Recognizing the inaccessibility of these truths, the Platonist seems rationally committed to suspending her mathematical beliefs, or else abandoning the Platonist commitments that generated the problem in the first place. What is it, though, that drives the Platonist to skepticism? Here is a natural answer: It would be irrational to retain beliefs that you come to think are not in any way explanatorily connected to their subject matter. In other words, an explanatory constraint on rational belief drives the access problem. Other answers have been offered. For instance, Hartry Field (1989) frames the problem as a challenge to explain the reliability of mathematical beliefs.

¹⁹ Though see Setiya (2012: 96) for a more nuanced construal.

But, as others have argued, it is trivially easy to meet this challenge (given the necessity of mathematical truths).²⁰ The access problem has no bite unless it's underwritten by something like [EXC]. But it *does* have bite. So something like [EXC] must be true.

4. The Missing Explanatory Connection

Having seen some motivations for [EXC], we are now ready to formulate an explanatory problem of induction:

[G1] There is no explanatory connection between our beliefs about the future and the future facts, and we recognize that this is the case.

[G2] If [G1], then we are not rationally entitled to retain our beliefs about the future.

[G3] We are not rationally entitled to retain our beliefs about the future.

Premiss [G2] is an immediate entailment of our explanatory constraint, [EXC]. As for [G1], unless we admit backwards causation or something similar, it is hard to see how future facts could explain our beliefs about the future. And except for a very restricted category of predictive inferences having to do with how we ourselves will act in the future, it is hard to see how our beliefs might explain any future facts. Specifically in the Thames and bismuth cases, we don't see how the fact that we believe that there will be a low tide tomorrow could explain the fact that there will be one or how the fact that we believe that a sample of bismuth will melt at a specific temperature tomorrow could explain the fact that it will melt at that temperature. If facts about the future neither explain nor are explained by our beliefs about the future, then there would seem to be no explanatory connection between them, just as [G1] says.

²⁰ See, e.g., Pust (2004), Baras (2017), and Clarke-Doane (forthcoming).

The material theory of induction is not equipped to answer this explanatory problem of induction. Consider again the predictive inference [T] that concludes with the proposition that tomorrow, there will be exactly one low tide on the Thames at London Bridge during daylight hours. Suppose we think inference [T] is reasonable, and having drawn it, we come to believe that tomorrow, there will be exactly one low tide on the Thames at London Bridge during daylight hours. Up to this point, our belief is rationally permissible (as far as [EXC] and [G2] are concerned). We are justified in believing that there will be exactly one low tide on the Thames tomorrow. When asked for our justification, we may respond just as Norton suggests, citing facts about estuary rivers and tides. But suppose we now consider the question of whether there is any explanatory connection between the fact of our belief and the fact believed. If, on reflection, we were to come to believe that there is no such connection, then we would have to give up our belief about tomorrow's tide. The justification that once secured our belief vanishes.

This highlights the main difference between the explanatory problem of induction and traditional formulations. Traditional formulations aim to show that beliefs based on predictive inference were never justified to begin with. The explanatory formulation aims to show that our ordinary sources of justification are undermined, or defeated, by explanatory revelations—which is entirely compatible with holding that the beliefs were justified prior to defeat. (By way of analogy, consider two sorts of skeptical attacks on some intuition-driven metaphysical principle. The first attack denies that intuition ever justified the principle in the first place. The second concedes that intuition is a source of prima facie justification, but insists that once we appreciate the unreliability of intuition, that justification is undermined.)

The aforementioned motivation for [G1]—namely, that the future tides don't cause our beliefs about future tides, and vice versa—is admittedly oversimplified. After all, there are other

ways for a pair of facts to be explanatorily connected. Even if the former don't explain the latter or vice versa, the idea goes, there may nevertheless be some further fact serving as a "third-factor" that explains both. Indeed, this is Goldman's (1967) strategy for accommodating inductive knowledge within a causal theory of knowledge, and it is a prominent strategy for resisting explanatory challenges in metaethics (e.g., Enoch 2010 and Wielenberg 2010). What might the third factor be that is supposed to explain both beliefs about the future and the future facts? We turn now to two strategies for resisting [G1] by identifying some such third factor.

The first candidate is [PI] itself, the fact that the future will resemble the past. That the future will be like the past (together with the fact that there have been regular tides in the past) explains the fact that there will be a low tide on the Thames tomorrow. And that the future will be like the past explains our belief that there will be a low tide tomorrow by way of natural selection. A creature that expects future states to resemble past states (in relevant respects) is bound to outperform one that lacks this expectation. And why is that? Because the future *does* resemble the past. In the memorable words of Quine (1969, 13): "Creatures inveterately wrong in their inductions have a pathetic but praiseworthy tendency to die before reproducing their kind."

Here is the problem with this line of reasoning (cf. Setiya 2012, 106-107). The putative fact that the future resembles the past is no part of the evolutionary explanation for why our inductive practices enhanced reproductive success for our ancestors. Rather, the explanation is that *their* future (a.k.a. the distant past) resembled *their* past (a.k.a. the somewhat more distant past). The facts about *our* future (a.k.a. the future) have no role to play in the evolutionary explanation; the evolutionary explanation features only facts about the past. Hence, the principle of induction cannot serve as a third factor explaining both our beliefs about the future and the future facts (since

it doesn't explain our beliefs). Therefore, the principle of induction cannot serve as grounds for denying [G1].

The second strategy we'll consider points to the laws of nature as the third factor explaining both our beliefs about the future and the future facts. According to this second strategy, the laws figure in the explanation of why the tides behave as they do and, in particular, why there will be exactly one low tide tomorrow. Additionally, the laws explain the regularities we have observed in the behavior of the tides, which in turn explain why we form the beliefs we do about future tides.²¹

This strategy is evidently available only to those who accept a "governing" conception of laws, according to which the laws are not merely regularities or descriptions thereof but rather some further thing underwriting and explaining those regularities. The strategy won't work in a nongoverning, Humean framework, since it won't be true that the tides behave as they do *because* the laws are as they are, and we lose the envisaged third-factor explanation. So for the sake of argument, let's assume that a governing account of the laws of nature is correct. We will argue that even if one accepts a governing account of laws, law-based third-factor explanations fail. As with the first strategy we considered, which appeals to the principle of induction, the core problem with the second strategy is that our beliefs are explained by the laws *up to now*, while the future facts are explained by the laws *in the future*. But before developing our argument, we want to briefly review a related criticism of a governing-law justification of induction.

Beebe (2011) considers two strategies for justifying induction by way of an appeal to laws of nature. Each strategy deploys the same basic inference pattern: use inference to the best explanation to infer the existence of a law of nature (which involves some kind of natural

²¹ Cf. Setiya (2012, 106-107) and Fairchild and Hawthorne (2018: 52).

necessity), and then deduce unobserved instances of the law. Beebee argues that the inference to the best explanation fails to deliver necessary connections in the future. We want to explain some observations that we have made. But from a “neutral” starting point—one where we neither believe nor doubt any proposition about the character of the laws—we have no reason to think that a timeless law, as opposed to a time-restricted law, does a better job explaining our observations. As Beebee puts it (510): “If *all* we are trying to explain is the fact that the observed *F*s have been *G*s, then the hypothesis that *F* and *G* have been necessarily connected *so far* is surely just as good a candidate explanation as the hypothesis that *F* and *G* are timelessly necessarily connected.”

The question of whether an explanation in terms of time-restricted laws is as good as or better than an explanation in terms of timeless laws is a difficult one. But we do not need to address it here because our objection to a governing-law approach to justifying predictive inference has a different starting point than Beebee’s and so must make use of different machinery. Beebee is considering an explanatory inference from a neutral starting point. But our explanatory constraint does not presuppose that we are starting from a neutral position. We need to consider the possibility that one is already committed to the existence of a specific law of nature, which one could then call on to provide the missing explanatory connection between one’s beliefs about the future and the future facts.²² We think that even in cases where one is initially justified in believing that there is a law of nature that entails both some future facts and one’s beliefs about those facts, the explanatory challenge is not satisfactorily met. To see how this could be the case, we need to follow a brief digression into the relation between explanation and coincidence.

²² Plausibly, when one sees that one had no good reason to infer a timeless law in the first place, but that a time-restricted law was preferable, one ought to give up believing that the timeless law exists. But we will not press this challenge further here.

David Faraci (2019), following Marc Lange (2010, 2016), argues that an explanation of some phenomenon shows it to be noncoincidental only if the explanation is sufficiently unified.

Consider the following example from Faraci (2019, 6):

It might seem quite the coincidence if you and [your colleague David] were in Minsk at the same time. We can offer the explanation: you are spending your sabbatical in Minsk, and [David is] on vacation in Minsk. Clearly, this does nothing to eliminate the coincidence.

There's some temptation to say this is no explanation at all. But even if the conjunctive fact that you're on sabbatical there and David is on vacation there explains why you're both in Minsk at the same time, you should not think the fact that you're both in Minsk is noncoincidental. The conjunctive fact does not provide a logically or epistemologically satisfactory account of the fact that you are both in Minsk. So, the conjunctive fact does not dispel the appearance of coincidence.²³

Faraci suggests (and we agree) that the conjunctive explanation does not dispel the appearance of coincidence because your presence and David's presence are being explained by different, separable conjuncts. The alleged explanation itself lacks the sort of unity required to properly bind together the fact that you are in Minsk and the fact that David is in Minsk. Faraci offers the following account of what it takes for an explanation of some phenomena to be unified:

UNIFIED. E is a unified explanation for the members of Γ if and only if no isolable part of E explains some members of Γ at least as well as E but fails to explain other members of Γ at least as well as E.²⁴

²³ We want to be careful here. It is relatively easy to confuse (1) a metaphysical question about whether and under what conditions some fact is a coincidence, (2) a logical or epistemological question about whether and under what conditions one ought to think that some fact is a coincidence, and (3) a psychological question about whether and under what conditions one actually thinks that some fact is a coincidence. In this paper, we are concerned with the logical or epistemological question.

²⁴ See his (2019: 6 n.17) for some possible complications.

The explanation *that David is on sabbatical in Minsk and you are on vacation in Minsk* is not a unified explanation for {your being in Minsk, David's being in Minsk}. That's because the explanation has an isolable part (viz. the first conjunct) that explains one member of the explanandum (viz. that you are in Minsk) at least as well as the whole conjunction but fails to explain another member of the explanandum (viz. that David is in Minsk).

The claim that disunified explanations shouldn't dispel the appearance of coincidence guides our understanding of our explanatory constraint [EXC]. Recall that [EXC] prohibits believers from acknowledging the absence of an "appropriate" explanatory connection between belief and what's believed. What's an "appropriate" explanatory connection? At the very least, it should be one that reveals the accuracy of the belief to be noncoincidental. For, as we saw in §3, the explanatory constraint is closely associated with a non-accidentality constraint: one doesn't know that *p* if it's an accident (or a coincidence) that one has correct beliefs regarding *p*, and continued belief in *p* is not rational if one recognizes that belief in *p* is at best accidental (or coincidental).²⁵

Thus, the envisaged third-factor explanation in terms of the laws of nature serves as a reason to reject [G1] *only if* it qualifies as a unified explanation. We now argue that the strategy fails: the laws of nature are not unified in the relevant way. As applied to the Thames inference, the explanandum would be {there will be a low tide tomorrow, you believe that there will be a low tide tomorrow}, and the explanation is that the laws of nature are thus-and-so. We claim that different parts of the explanation are responsible for different members of the explanandum. As with the Quinean response, what explains (and explains just as well) your *beliefs* about future tides

²⁵ See Korman and Locke (forthcoming: §8) for further reason to think that the explanatory constraint on rational belief is not satisfied by such third-factor explanations.

is that the laws of nature have been thus-and-so up to now, whereas what explains the *facts* about future tides is that the laws of nature are whatever they are in the future.

Do these count as “isolable parts” of the claim that the laws of nature are thus-and-so, in the sense required by UNIFIED? They aren’t conjuncts of the proffered explanation, so the application of UNIFIED to this case admittedly isn’t as straightforward as it is in Faraci’s Minsk example. Still, there is independent reason to adopt a liberal understanding of “isolable parts.” Consider the following imaginary version of what Turkheimer (2016) calls *weak genetic explanation*:

Martia is a behavioral geneticist studying divorce. Last year, she collected data on divorce and genetics in seven European countries. All of the people in the seven countries Martia studied were genetically very similar to one another, and the divorce rate was the same in each country. Moreover, Martia found that in each of her seven samples, there was a statistically significant association between genes and divorce rate. Martia initially thought that there might be genes for divorce. But when she looked more closely at her data, she found that the specific genes that were significant in her model were entirely different for each country.

That all of the people in the seven countries were genetically very similar to one another is a (weak) genetic explanation of the divorce rates. And yet, it’s a coincidence that the divorce rate is the same in each country. UNIFIED should capture why this is a coincidence, but in order to do so, we need to understand the specific, local genetic story for each country as an isolable part of the claim that all of the people in the seven countries were genetically very similar to one another. The need for a more flexible understanding of “isolable part” then clears the way for treating *it’s a law that thus-and-so* as decomposing into two isolable parts: that the laws have been thus-and-so and that the laws will be thus-and-so.

For these reasons, we contend that the envisaged third-factor explanation is insufficiently unified to render the connection between our beliefs about the future and the future facts

noncoincidental, and therefore, appealing to the laws fails to serve as the sort of “appropriate” explanatory connection required by [EXC].

5. Concluding Remarks

Norton’s material theory of induction likely has the resources to address the problem of induction as traditionally understood. But we think that the material theory is not equipped to handle the explanatory problem of induction. Norton may well insist that the material theory was never meant to handle *every* problem of induction. Indeed, he is explicit that his goal is to address “a particular, brief demonstration of the impossibility of justifying any rule of inductive inference,” insisting that separate problems of induction must be taken one at a time if there is to be any progress on the matter, and lamenting that “too often, discussions of the problem bleed off into other problems attached to inductive inference ... and then further into a vague foreboding that inductive inference is philosophically problematic in all its aspects” (2014: 674). We are sympathetic, and we don’t mean to suggest that it is a defect of the material theory that it is not equipped to block the explanatory problem.

Still, friends of predictive inference need to find *something* to say in response to the explanatory problem. We’ve seen that there are powerful reasons for thinking that there is some sort of explanatory constraint on justification—not to mention that the explanatory constraint that we have developed has clear, empiricist-friendly motivations of the sort that we think will be appealing to Norton and many other philosophers of science. Is there some way of amending the explanatory constraint, so that it retains its core virtues without leading to skepticism about predictive inference? We hope so, but it isn’t obvious how such an amendment might go. A flatfooted strategy would be to simply add a clause licensing inductive inferences:

[EXC'] If *S* believes that there is no appropriate explanatory connection between her belief that *p* and the fact that *p* and *S* has no beliefs $q_1...q_n$ from which *p* follows inductively, then *S* is not rationally entitled to believe that *p*.

But we find this sort of amendment unilluminating and seemingly ad hoc. What's so special about beliefs formed by induction, that they are exempt from the explanatory constraint?

One might suggest that predictive inferences aren't special; they're simply one case among many in which one believes some proposition *p* on the basis of beliefs that are themselves explanatorily connected to the associated facts. The belief that there will be a tide on the Thames tomorrow is based on the belief that there have been tides on all previous days, and while the former belief is not explanatorily connected to the associated future tidal facts, the latter belief is explanatorily connected to the associated past tidal facts and supports the former belief. This suggests a somewhat different amendment:

[EXC''] If *S* believes that there is no appropriate explanatory connection between her belief that *p* and the fact that *p* and *S* believes that she has no beliefs $q_1...q_n$ that (i) support *p* and (ii) are explained by facts $q_1...q_n$, then *S* is not rationally entitled to believe that *p*.²⁶

Since we do take there to be beliefs about past tides that explain our beliefs about past tides and that (inductively) support our beliefs about future tides, [EXC''] won't prescribe withholding belief about future tides.

But [EXC''] is problematic as well, since it threatens to make it trivially easy to avoid the access problems discussed in Section 3. For instance, take some utterly mundane reason to think that it's wrong to set fire to cats, for instance that it's quite painful for them. By the lights of [EXC''], even if you concede that there is no explanatory connection whatsoever between your moral beliefs and the moral facts, your belief that it's wrong to set fire to cats is in good standing so long as you accept that your beliefs about what's painful for cats are explained by facts about

²⁶ See Korman (2019: 356) for a proposal along these lines.

what is painful for them. One might respond that whatever support beliefs about pain ordinarily supply for beliefs about wrongness is undercut by the explanatory concession that one has no access to moral facts. And rightly so! But [EXC"] is powerless to explain why it's undercut.

We're not sure whether to be optimistic or pessimistic about the prospects of finding an induction-friendly explanatory constraint that can do the work that an explanatory constraint needs to do. We leave it as an open problem to find one.

References

- Achinstein, P. (2008). Evidence. *The Routledge Companion to Philosophy of Science*, 337-348.
- Achinstein, P. (2010). The War on Induction: Whewell Takes On Newton and Mill (Norton Takes On Everyone). *Philosophy of Science*, 77(5), 728-739.
- Baras, D. (2017). Our reliability is in principle explainable. *Episteme*, 14(2), 197-211.
- Beebe, H. (2011). Necessary Connections and the Problem of Induction. *Noûs* 45(3), 504-527
- Benacerraf, P. (1973) Mathematical Truth. *Journal of Philosophy* 70(19), 661-679.
- Carnap, R. (1950). *Logical Foundations of Probability*. Chicago: University of Chicago Press.
- Clarke-Doane, J. (forthcoming), What is the Benacerraf Problem? In Fabrice Pataut (ed.), *New Perspectives on the Philosophy of Paul Benacerraf*.
- Dogramaci, S. (2017). Explaining our Moral Reliability. *Pacific Phil Quarterly* 98, 71-86.
- Fairchild, M., & Hawthorne, J. (2018). Against conservatism in metaphysics. *Royal Institute of Philosophy Supplements*, 82, 45-75.
- Faraci, D. (2019), Groundwork for an Explanationist Account of Epistemic Coincidence, *Philosophers' Imprint* 19(4), 1-26.
- Field, H. (1989), *Realism, Mathematics, and Modality* (Oxford: Blackwell).
- Goldman, A. I. (1967). A causal theory of knowing. *The journal of Philosophy*, 64(12), 357-372.
- Goldman, A. I., & McLaughlin, B. P. (Eds.). (2019). *Metaphysics and Cognitive Science*. Oxford University Press.
- Haack, S. (1993). *Evidence and Inquiry: Towards Reconstruction in Epistemology*. Oxford: Blackwell.
- Horowitz, S. (2014). Epistemic Akrasia. *Noûs* 48(4), 718-744.
- Huemer, M. (2001), *Skepticism and the Veil of Perception* (Lanham: Rowman & Littlefield).
- Jacobson, A. (1987). The Problem of Induction: What is Hume's Argument? *Pacific Philosophical Quarterly* 68, 265-284.
- Joyce, R. (2006), *The Evolution of Morality* (Cambridge: MIT).

- Kelly, T. (2010). Hume, Norton, and Induction without Rules. *Philosophy of Science*, 77(5), 754-764.
- Korman, D. Z. (2019). Debunking Arguments in Metaethics and Metaphysics. *Metaphysics and Cognitive Science*, 337-363.
- Korman, D. Z. and Locke, D. (forthcoming). Against Minimalist Responses to Moral Debunking Arguments. *Oxford Studies in Metaethics*.
- Lange, M. (2010). What are Mathematical Coincidences? *Mind* 119(474), 307-340.
- Lange, M. (2016). *Because without Cause: Non-Causal Explanations in Science and Mathematics* (Oxford: Oxford University Press).
- Livengood, J., & Edwards, A. (forthcoming). Calibrating Chromatography: How Tswett Broke the Experimenters' Regress. *British Journal for the Philosophy of Science*.
- Norton, J. D. (2003). A material theory of induction. *Philosophy of Science*, 70(4), 647-670.
- Norton, J. D. (2010). There are no universal rules for induction. *Philosophy of Science*, 77(5), 765-777.
- Norton, J. D. (2011). History of science and the material theory of induction: Einstein's quanta, mercury's perihelion. *European Journal for Philosophy of Science*, 1(1), 3-27.
- Norton, J. D. (2014). A material dissolution of the problem of induction. *Synthese*, 191(4), 671-690.
- Norton, J. (Forthcoming). *The Material Theory of Induction*. Manuscript available at URL=< https://www.pitt.edu/~jdnorton/homepage/cv.html#material_theory>
- Okasha, S. (2005). Does Hume's Argument against Induction Rest on a Quantifier-Shift Fallacy? *Proceedings of the Aristotelian Society* 105, 237-255.
- Popper, K. (1972). *Objective Knowledge*. Oxford: Oxford University Press.
- Pritchard, D. (2005). *Epistemic Luck*. Oxford: Oxford University Press.
- Pritchard, D. (2007). Anti-luck epistemology. *Synthese*, 158(3), 277-297.
- Pritchard, D. (2009). Safety-based epistemology: Whither now? *Journal of Philosophical Research*, 34, 33-45.
- Psillos, S. and Curd, M. (2008). *The Routledge Companion to Philosophy of Science*. New York: Routledge.

- Pust, J. (2004). On Explaining Knowledge of Necessity. *Dialectica* 58: 71-87.
- Quine, W. (1969). Natural Kinds. *Essays in Honor of Carl G. Hempel*, 5-23.
- Salmon, W. (1981). Rational Prediction. *The British Journal for the Philosophy of Science* 32(2), 115-125.
- Setiya, K. (2012), *Knowing Right from Wrong*. Oxford: Oxford University Press.
- Steel, D. (2005). The facts of the matter: A discussion of Norton's material theory of induction. *Philosophy of Science*, 72(1), 188-197.
- Street, S. (2006). A Darwinian dilemma for realist theories of value. *Philosophical Studies*, 127(1), 109-166.
- Turkheimer, E. (2016). Weak genetic explanation 20 years later: reply to Plomin et al.(2016). *Perspectives on Psychological Science*, 11(1), 24-28.
- Unger, P. (1967). Experience and factual knowledge. *The Journal of Philosophy*, 64(5), 152-173.
- Unger, P. (1968). An analysis of factual knowledge. *The Journal of Philosophy*, 65(6), 157-170.
- White, R. (2010). You Just Believe That Because. *Philosophical Perspectives* 24: 573-615.
- Worrall, J. (Ed.). (2010). For universal rules, against induction. *Philosophy of Science*, 77(5), 740-753.