

The Problem of Induction

First published Wed Mar 21, 2018; substantive revision Tue Nov 22, 2022

We generally think that the observations we make are able to justify some expectations or predictions about observations we have not yet made, as well as general claims that go beyond the observed. For example, the observation that bread of a certain appearance has thus far been nourishing seems to justify the expectation that the next similar piece of bread I eat will also be nourishing, as well as the claim that bread of this sort is generally nourishing. Such inferences from the observed to the unobserved, or to general laws, are known as “inductive inferences”.

The original source of what has become known as the “problem of induction” is in Book 1, part iii, section 6 of *A Treatise of Human Nature* by David Hume, published in 1739 (Hume 1739). In 1748, Hume gave a shorter version of the argument in Section iv of *An enquiry concerning human understanding* (Hume 1748). Throughout this article we will give references to the *Treatise* as “T”, and the *Enquiry* as “E”.

Hume asks on what grounds we come to our beliefs about the unobserved on the basis of inductive inferences. He presents an argument in the form of a dilemma which appears to rule out the possibility of any reasoning from the premises to the conclusion of an inductive inference. There are, he says, two possible types of arguments, “demonstrative” and “probable”, but neither will serve. A demonstrative argument produces the wrong kind of conclusion, and a probable argument would be circular. Therefore, for Hume, the problem remains of how to explain why we form any conclusions that go beyond the past instances of which we have had experience (T. 1.3.6.10). Hume stresses that he is not disputing that we do draw such inferences. The challenge, as he sees it, is to understand the “foundation” of the inference—the “logic” or “process of argument” that it is based upon (E. 4.2.21). The problem of meeting this challenge, while evading Hume’s argument against the possibility of doing so, has become known as “the problem of induction”.

Hume’s argument is one of the most famous in philosophy. A number of philosophers have attempted solutions to the problem, but a significant number have embraced his conclusion that it is insoluble. There is also a wide spectrum of opinion on the significance of the problem. Some have argued that Hume’s argument does not establish any far-reaching skeptical conclusion, either because it was never intended to, or because the argument is in some way misformulated. Yet many have regarded it as one of the most profound philosophical challenges imaginable since it seems to call into question the justification of one of the most fundamental ways in which we form knowledge. Bertrand Russell, for example, expressed the view that if Hume’s problem cannot be solved, “there is no intellectual difference between sanity and insanity” (Russell 1946: 699).

In this article, we will first examine Hume’s own argument, provide a reconstruction of it, and then survey different responses to the problem which it poses.

- 1. Hume’s Problem
- 2. Reconstruction
- 3. Tackling the First Horn of Hume’s Dilemma
 - 3.1 Synthetic *a priori*
 - 3.2 The Nomological-Explanatory solution

- 3.3 Bayesian solution
 - 3.4 Partial solutions
 - 3.5 The combinatorial approach
 - 4. Tackling the Second Horn of Hume’s Dilemma
 - 4.1 Inductive Justifications of Induction
 - 4.2 No Rules
 - 5. Alternative Conceptions of Justification
 - 5.1 Postulates and Hinges
 - 5.2 Ordinary Language Dissolution
 - 5.3 Pragmatic vindication of induction
 - 5.4 Formal Learning Theory
 - 5.5 Meta-induction
 - 6. Living with Inductive Skepticism
 - Bibliography
 - Academic Tools
 - Other Internet Resources
 - Related Entries
-

1. Hume’s Problem

Hume introduces the problem of induction as part of an analysis of the notions of cause and effect. Hume worked with a picture, widespread in the early modern period, in which the mind was populated with mental entities called “ideas”. Hume thought that ultimately all our ideas could be traced back to the “impressions” of sense experience. In the simplest case, an idea enters the mind by being “copied” from the corresponding impression (T. 1.1.1.7/4). More complex ideas are then created by the combination of simple ideas (E. 2.5/19). Hume took there to be a number of relations between ideas, including the relation of causation (E. 3.2). (For more on Hume’s philosophy in general, see Morris & Brown 2014).

For Hume, the relation of causation is the only relation by means of which “we can go beyond the evidence of our memory and senses” (E. 4.1.4, T. 1.3.2.3/74). Suppose we have an object present to our senses: say gunpowder. We may then infer to an effect of that object: say, the explosion. The causal relation links our past and present experience to our expectations about the future (E. 4.1.4/26).

Hume argues that we cannot make a causal inference by purely *a priori* means (E. 4.1.7). Rather, he claims, it is based on experience, and specifically experience of constant conjunction. We infer that the gunpowder will explode on the basis of past experience of an association between gunpowder and explosions.

Hume wants to know more about the basis for this kind of inference. If such an inference is made by a “chain of reasoning” (E. 4.2.16), he says, he would like to know what that reasoning is. In general, he claims that the inferences depend on a transition of the form:

I have found that such an object has always been attended with such an effect, and I foresee, that other objects, which are, in appearance, similar, will be attended with similar effects. (E. 4.2.16)

In the *Treatise*, Hume says that

if Reason determin'd us, it would proceed upon that principle *that instances, of which we have had no experience, must resemble those, of which we have had experience, and that the course of nature continues always uniformly the same.* (T. 1.3.6.4)

For convenience, we will refer to this claim of similarity or resemblance between observed and unobserved regularities as the “Uniformity Principle (UP)”. Sometimes it is also called the “Resemblance Principle”, or the “Principle of Uniformity of Nature”.

Hume then presents his famous argument to the conclusion that there can be no reasoning behind this principle. The argument takes the form of a dilemma. Hume makes a distinction between relations of ideas and matters of fact. Relations of ideas include geometric, algebraic and arithmetic propositions, “and, in short, every affirmation, which is either intuitively or demonstratively certain”. “Matters of fact”, on the other hand are empirical propositions which can readily be conceived to be other than they are. Hume says that

All reasonings may be divided into two kinds, namely, demonstrative reasoning, or that concerning relations of ideas, and moral reasoning, or that concerning matter of fact and existence. (E. 4.2.18)

Hume considers the possibility of each of these types of reasoning in turn, and in each case argues that it is impossible for it to supply an argument for the Uniformity Principle.

First, Hume argues that the reasoning cannot be demonstrative, because demonstrative reasoning only establishes conclusions which cannot be conceived to be false. And, he says,

it implies no contradiction that the course of nature may change, and that an object seemingly like those which we have experienced, may be attended with different or contrary effects. (E. 4.2.18)

It is possible, he says, to clearly and distinctly conceive of a situation where the unobserved case does not follow the regularity so far observed (E. 4.2.18, T. 1.3.6.5/89).

Second, Hume argues that the reasoning also cannot be “such as regard matter of fact and real existence”. He also calls this “probable” reasoning. All such reasoning, he claims, “proceed upon the supposition, that the future will be conformable to the past”, in other words on the Uniformity Principle (E. 4.2.19).

Therefore, if the chain of reasoning is based on an argument of this kind it will again be relying on this supposition, “and taking that for granted, which is the very point in question”. (E. 4.2.19, see also T. 1.3.6.7/90). The second type of reasoning then fails to provide a chain of reasoning which is not circular.

In the *Treatise* version, Hume concludes

Thus, not only our reason fails us in the discovery of the *ultimate connexion* of causes and effects, but even after experience has inform'd us of their *constant conjunction*, 'tis impossible for us to satisfy ourselves by our reason, why we shou'd extend that experience beyond those particular instances, which have fallen under our observation. (T. 1.3.6.11/91–2)

The conclusion then is that our tendency to project past regularities into the future is not underpinned by reason. The problem of induction is to find a way to avoid this conclusion, despite Hume's argument.

After presenting the problem, Hume does present his own “solution” to the doubts he has raised (E. 5, T. 1.3.7–16). This consists of an explanation of what the inductive inferences are driven by, if not reason. In the *Treatise* Hume raises the problem of induction in an explicitly contrastive way. He asks whether the transition involved in the inference is produced

by means of the understanding or imagination; whether we are determin'd by reason to make the transition, or by a certain association and relation of perceptions? (T. 1.3.6.4)

And he goes on to summarize the conclusion by saying

When the mind, therefore, passes from the idea or impression of one object to the idea or belief of another, it is not determin'd by reason, but by certain principles, which associate together the ideas of these objects, and unite them in the imagination. (T. 1.3.6.12)

Thus, it is the imagination which is taken to be responsible for underpinning the inductive inference, rather than reason.

In the *Enquiry*, Hume suggests that the step taken by the mind,

which is not supported by any argument, or process of the understanding ... must be induced by some other principle of equal weight and authority. (E. 5.1.2)

That principle is “custom” or “habit”. The idea is that if one has seen similar objects or events constantly conjoined, then the mind is inclined to expect a similar regularity to hold in the future. The tendency or “propensity” to draw such inferences, is the effect of custom:

... having found, in many instances, that any two kinds of objects, flame and heat, snow and cold, have always been conjoined together; if flame or snow be presented anew to the senses, the mind is carried by custom to expect heat or cold, and to *believe*, that such a quality does exist and will discover itself upon a nearer approach. This belief is the necessary result of placing the mind in such circumstances. It is an operation of the soul, when we are so situated, as unavoidable as to feel the passion of love, when we receive benefits; or hatred, when we meet with injuries. All these operations are a species of natural instincts, which no reasoning or process of the thought and understanding is able, either to produce, or to prevent. (E. 5.1.8)

Hume argues that the fact that these inferences do follow the course of nature is a kind of “pre-established harmony” (E. 5.2.21). It is a kind of natural instinct, which may in fact be more effective in making us successful in the world, than if we relied on reason to make these inferences.

2. Reconstruction

Hume’s argument has been presented and formulated in many different versions. There is also an ongoing lively discussion over the historical interpretation of what Hume himself intended by the argument. It is therefore difficult to provide an unequivocal and uncontroversial reconstruction of Hume’s argument. Nonetheless, for the purposes of organizing the different responses to Hume’s problem that will be discussed in this article, the following reconstruction will serve as a useful starting point.

Hume’s argument concerns specific inductive inferences such as:

All observed instances of *A* have been *B*.

The next instance of *A* will be *B*.

Let us call this “inference *I*”. Inferences which fall under this type of schema are now often referred to as cases of “simple enumerative induction”.

Hume’s own example is:

All observed instances of bread (of a particular appearance) have been nourishing.

The next instance of bread (of that appearance) will be nourishing.

Hume’s argument then proceeds as follows (premises are labeled as P, and subconclusions and conclusions as C):

P1. There are only two kinds of arguments: demonstrative and probable (Hume’s fork).

P2. Inference *I* presupposes the Uniformity Principle (UP).

1st horn:

P3. A demonstrative argument establishes a conclusion whose negation is a contradiction.

P4. The negation of the UP is not a contradiction.

C1. There is no demonstrative argument for the UP (by P3 and P4).

2nd horn:

P5. Any probable argument for UP presupposes UP.

P6. An argument for a principle may not presuppose the same principle (Non-circularity).

C2. There is no probable argument for the UP (by P5 and P6).

Consequences:

C3. There is no argument for the UP (by P1, C1 and C2).

P7. If there is no argument for the UP, there is no chain of reasoning from the premises to the conclusion of any inference that presupposes the UP.

C4. There is no chain of reasoning from the premises to the conclusion of inference *I* (by P2, C3 and P7).

P8. If there is no chain of reasoning from the premises to the conclusion of inference *I*, the inference is not justified.

C5. Inference *I* is not justified (by C4 and P8).

There have been different interpretations of what Hume means by “demonstrative” and “probable” arguments. Sometimes “demonstrative” is equated with “deductive”, and probable with “inductive” (e.g., Salmon 1966). Then the first horn of Hume’s dilemma would eliminate the possibility of a deductive argument, and the second would eliminate the possibility of an inductive argument. However, under this interpretation, [premise P3](#) would not hold, because it is possible for the conclusion of a deductive argument to be a non-necessary proposition. Premise [P3](#) could be modified to say that a demonstrative (deductive) argument establishes a conclusion that cannot be false if the premises are true. But then it becomes possible that the supposition that the future resembles the past, which is not a necessary proposition, could be established by a deductive argument from some premises, though not from *a priori* premises (in contradiction to conclusion [C1](#)).

Another common reading is to equate “demonstrative” with “deductively valid with *a priori* premises”, and “probable” with “having an empirical premise” (e.g., Okasha 2001). This may be closer to the mark, if one thinks, as Hume seems to have done, that premises which can be known *a priori* cannot be false, and hence are necessary. If the inference is deductively valid, then the conclusion of the inference from *a priori* premises must also be necessary. What the first horn of the dilemma then rules out is the possibility of a deductively valid argument with *a priori* premises, and the second horn rules out any argument (deductive or non-deductive), which relies on an empirical premise.

However, recent commentators have argued that in the historical context that Hume was situated in, the distinction he draws between demonstrative and probable arguments has little to do with whether or not the argument has a deductive form (Owen 1999; Garrett 2002). In addition, the class of inferences that establish conclusions whose negation is a contradiction may include not just deductively valid inferences from *a priori* premises, but any inferences that can be drawn using *a priori* reasoning (that is, reasoning where the transition from premises to the conclusion makes no appeal to what we learn from observations). It looks as though Hume does intend the argument of the first horn to rule out any *a priori* reasoning, since he says that a change in the course of nature cannot be ruled out “by any demonstrative argument or abstract reasoning *a priori*” (E. 5.2.18). On this understanding, *a priori* arguments would be ruled out by the first horn of Hume’s dilemma, and empirical arguments by the second horn. This is the interpretation that I will adopt for the purposes of this article.

In Hume’s argument, the UP plays a central role. As we will see in [section 4.2](#), various authors have been doubtful about this principle. Versions of Hume’s argument have also been formulated which do not make reference to the UP. Rather they directly address the question of what arguments can be given in support of the transition from the premises to the conclusion of the specific inductive inference *I*. What arguments could lead us, for example, to infer that the next piece of bread will nourish from the observations of nourishing bread made so far? For the first horn of the argument, Hume’s argument can be directly applied. A demonstrative argument establishes a conclusion whose negation is a contradiction. The negation of the conclusion of the inductive inference is not a contradiction. It is not a contradiction that the next piece of bread is not nourishing. Therefore, there is no demonstrative argument for the conclusion of the inductive inference. In the second horn of the argument, the problem Hume raises is a circularity. Even if Hume is wrong that all inductive inferences depend on the UP, there may still be a circularity problem, but as we shall see in [section 4.1](#), the exact nature of the circularity needs to be carefully considered. But the main point at present is that the Humean argument is often formulated without invoking the UP.

Since Hume’s argument is a dilemma, there are two main ways to resist it. The first is to tackle the first horn and to argue that there is after all a demonstrative argument—here taken to mean an argument based on *a priori* reasoning—that can justify the inductive inference. The second is to tackle the second horn and to argue that there is after all a probable (or empirical) argument that can justify the inductive inference. We discuss the different variants of these two approaches in [sections 3](#) and [4](#).

There are also those who dispute the consequences of the dilemma. For example, some scholars have denied that Hume should be read as invoking a premise such [premise P8](#) at all. The reason, they claim, is that he was not aiming for an explicitly normative conclusion about justification such as [C5](#). Hume certainly is seeking a “chain of reasoning” from the premises of the inductive inference to the conclusion, and he thinks that an argument for the UP is necessary to complete the chain. However, one could think that there is no further premise regarding justification, and so the conclusion of his argument is simply [C4](#): there is no chain of reasoning from the premises to the conclusion of an inductive inference. Hume could then be, as Don Garrett and David Owen have argued, advancing a “thesis in cognitive psychology”, rather than making a normative claim about justification (Owen 1999; Garrett 2002). The

thesis is about the nature of the cognitive process underlying the inference. According to Garrett, the main upshot of Hume's argument is that there can be no reasoning process that establishes the UP. For Owen, the message is that the inference is not drawn through a chain of ideas connected by mediating links, as would be characteristic of the faculty of reason.

There are also interpreters who have argued that Hume is merely trying to exclude a specific kind of justification of induction, based on a conception of reason predominant among rationalists of his time, rather than a justification in general (Beauchamp & Rosenberg 1981; Baier 2009). In particular, it has been claimed that it is "an attempt to refute the rationalist belief that at least some inductive arguments are demonstrative" (Beauchamp & Rosenberg 1981: xviii). Under this interpretation, [premise P8](#) should be modified to read something like:

- If there is no chain of reasoning based on demonstrative arguments from the premises to the conclusion of inference *I*, then inference *I* is not justified.

Such interpretations do however struggle with the fact that Hume's argument is explicitly a two-pronged attack, which concerns not just demonstrative arguments, but also probable arguments.

The question of how expansive a normative conclusion to attribute to Hume is a complex one. It depends in part on the interpretation of Hume's own solution to his problem. As we saw in [section 1](#), Hume attributes the basis of inductive inference to principles of the imagination in the *Treatise*, and in the *Enquiry* to "custom", "habit", conceived as a kind of natural instinct. The question is then whether this alternative provides any kind of justification for the inference, even if not one based on reason. On the face of it, it looks as though Hume is suggesting that inductive inferences proceed on an entirely arational basis. He clearly does not think that they do not succeed in producing good outcomes. In fact, Hume even suggests that this operation of the mind may even be less "liable to error and mistake" than if it were entrusted to "the fallacious deductions of our reason, which is slow in its operations" (E. 5.2.22). It is also not clear that he sees the workings of the imagination as completely devoid of rationality. For one thing, Hume talks about the imagination as governed by *principles*. Later in the *Treatise*, he even gives "rules" and "logic" for characterizing what should count as a good causal inference (T. 1.3.15). He also clearly sees it as possible to distinguish between better forms of such "reasoning", as he continues to call it. Thus, there may be grounds to argue that Hume was not trying to argue that inductive inferences have no rational foundation whatsoever, but merely that they do not have the specific type of rational foundation which is rooted in the faculty of Reason.

All this indicates that there is room for debate over the intended scope of Hume's own conclusion. And thus there is also room for debate over exactly what form a premise (such as [premise P8](#)) that connects the rest of his argument to a normative conclusion should take. No matter who is right about this however, the fact remains that Hume has throughout history been predominantly read as presenting an argument for inductive skepticism.

There are a number of approaches which effectively, if not explicitly, take issue with [premise P8](#) and argue that providing a chain of reasoning from the premises to the conclusion is not a necessary condition for justification of an inductive inference. According to this type of approach, one may admit that Hume has shown that inductive inferences are not justified in the sense that we have reasons to think their conclusions true, but still think that weaker kinds of justification of induction are possible ([section 5](#)). Finally, there are some philosophers who do accept the skeptical conclusion [C5](#) and attempt to accommodate it. For example, there have been attempts to argue that inductive inference is not as central to scientific inquiry as is often thought ([section 6](#)).

3. Tackling the First Horn of Hume's Dilemma

The first horn of Hume's argument, as formulated above, is aimed at establishing that there is no demonstrative argument for the UP. There are several ways people have attempted to show that the first horn does not definitively preclude a demonstrative or *a priori* argument for inductive inferences. One possible escape route from the first horn is to deny [premise P3](#), which amounts to admitting the possibility of synthetic *a priori* propositions ([section 3.1](#)). Another possibility is to attempt to provide an *a priori* argument that the conclusion of the inference is probable, though not certain. The first horn of Hume's dilemma implies that there cannot be a demonstrative argument to the conclusion of an inductive inference because it is possible to conceive of the negation of the conclusion. For instance, it is quite possible to imagine that the next piece of bread I eat will poison me rather than nourish me. However, this does not rule out the possibility of a demonstrative argument that establishes only that the bread is highly likely to nourish, not that it definitely will. One might then also challenge [premise P8](#), by saying that it is not necessary for justification of an inductive inference to have a chain of reasoning from its premises to its conclusion. Rather it would suffice if we had an argument from the premises to the claim that the conclusion is probable or likely. Then an *a priori* justification of the inductive inference would have been provided. There have been attempts to provide *a priori* justifications for inductive inference based on Inference to the Best Explanation ([section 3.2](#)). There are also attempts to find an *a priori* solution based on probabilistic formulations of inductive inference, though many now think that a purely *a priori* argument cannot be found because there are empirical assumptions involved ([sections 3.3-3.5](#)).

3.1 Synthetic *a priori*

As we have seen in [section 1](#), Hume takes demonstrative arguments to have conclusions which are "relations of ideas", whereas "probable" or "moral" arguments have conclusions which are "matters of fact". Hume's distinction between "relations of ideas" and "matters of fact" anticipates the distinction drawn by Kant between "analytic" and "synthetic" propositions (Kant 1781). A classic example of an analytic proposition is "Bachelors are unmarried men", and a synthetic proposition is "My bike tyre is flat". For Hume, demonstrative arguments, which are based on *a priori* reasoning, can establish only relations of ideas, or analytic propositions. The association between a prioricity and analyticity underpins [premise P3](#), which states that a demonstrative argument establishes a conclusion whose negation is a contradiction.

One possible response to Hume's problem is to deny [premise P3](#), by allowing the possibility that *a priori* reasoning could give rise to synthetic propositions. Kant famously argued in response to Hume that such synthetic *a priori* knowledge is possible (Kant 1781, 1783). He does this by a kind of reversal of the empiricist programme espoused by Hume. Whereas Hume tried to understand how the concept of a causal or necessary connection could be based on experience, Kant argued instead that experience only comes about through the concepts or "categories" of the understanding. On his view, one can gain *a priori* knowledge of these concepts, including the concept of causation, by a transcendental argument concerning the necessary preconditions of experience. A more detailed account of Kant's response to Hume can be found in de Pierris and Friedman 2013.

3.2 The Nomological-Explanatory solution

The "Nomological-explanatory" solution, which has been put forward by Armstrong, Bonjour and Foster (Armstrong 1983; Bonjour 1998; Foster 2004) appeals to the principle of Inference to the Best Explanation (IBE). According to IBE, we should infer that the hypothesis which provides the best explanation of the evidence is probably true. Proponents of the Nomological-Explanatory approach take

Inference to the Best Explanation to be a mode of inference which is distinct from the type of “extrapolative” inductive inference that Hume was trying to justify. They also regard it as a type of inference which although non-deductive, is justified *a priori*. For example, Armstrong says “To infer to the best explanation is part of what it is to be rational. If that is not rational, what is?” (Armstrong 1983: 59).

The *a priori* justification is taken to proceed in two steps. First, it is argued that we should recognize that certain observed regularities require an explanation in terms of some underlying law. For example, if a coin persistently lands heads on repeated tosses, then it becomes increasingly implausible that this occurred just because of “chance”. Rather, we should infer to the better explanation that the coin has a certain bias. Saying that the coin lands heads not only for the observed cases, but also for the unobserved cases, does not provide an explanation of the observed regularity. Thus, mere Humean constant conjunction is not sufficient. What is needed for an explanation is a “non-Humean, metaphysically robust conception of objective regularity” (BonJour 1998), which is thought of as involving actual natural necessity (Armstrong 1983; Foster 2004).

Once it has been established that there must be some metaphysically robust explanation of the observed regularity, the second step is to argue that out of all possible metaphysically robust explanations, the “straight” inductive explanation is the best one, where the straight explanation extrapolates the observed frequency to the wider population. For example, given that a coin has some objective chance of landing heads, the best explanation of the fact that m/n heads have been so far observed, is that the objective chance of the coin landing heads is m/n . And this objective chance determines what happens not only in observed cases but also in unobserved cases.

The Nomological-Explanatory solution relies on taking IBE as a rational, *a priori* form of inference which is distinct from inductive inferences like inference *I*. However, one might alternatively view inductive inferences as a special case of IBE (Harman 1968), or take IBE to be merely an alternative way of characterizing inductive inference (Henderson 2014). If either of these views is right, IBE does not have the necessary independence from inductive inference to provide a non-circular justification of it.

One may also object to the Nomological-Explanatory approach on the grounds that regularities do not necessarily require an explanation in terms of necessary connections or robust metaphysical laws. The viability of the approach also depends on the tenability of a non-Humean conception of laws. There have been several serious attempts to develop such an account (Armstrong 1983; Tooley 1977; Dretske 1977), but also much criticism (see J. Carroll 2016).

Another critical objection is that the Nomological-Explanatory solution simply begs the question, even if it is taken to be legitimate to make use of IBE in the justification of induction. In the first step of the argument we infer to a law or regularity which extends beyond the spatio-temporal region in which observations have been thus far made, in order to predict what will happen in the future. But why could a law that only applies to the observed spatio-temporal region not be an equally good explanation? The main reply seems to be that we can see *a priori* that laws with temporal or spatial restrictions would be less good explanations. Foster argues that the reason is that this would introduce more mysteries:

For it seems to me that a law whose scope is restricted to some particular period is more mysterious, inherently more puzzling, than one which is temporally universal. (Foster 2004)

3.3 Bayesian solution

Another way in which one can try to construct an *a priori* argument that the premises of an inductive inference make its conclusion probable, is to make use of the formalism of probability theory itself. At the time Hume wrote, probabilities were used to analyze games of chance. And in general, they were used to address the problem of what we would expect to see, given that a certain cause was known to be operative. This is the so-called problem of “direct inference”. However, the problem of induction concerns the “inverse” problem of determining the cause or general hypothesis, given particular observations.

One of the first and most important methods for tackling the “inverse” problem using probabilities was developed by Thomas Bayes. Bayes’s essay containing the main results was published after his death in 1764 (Bayes 1764). However, it is possible that the work was done significantly earlier and was in fact written in direct response to the publication of Hume’s *Enquiry* in 1748 (see Zabell 1989: 290–93, for discussion of what is known about the history).

We will illustrate the Bayesian method using the problem of drawing balls from an urn. Suppose that we have an urn which contains white and black balls in an unknown proportion. We draw a sample of balls from the urn by removing a ball, noting its color, and then putting it back before drawing again.

Consider first the problem of direct inference. Given the proportion of white balls in the urn, what is the probability of various outcomes for a sample of observations of a given size? Suppose the proportion of white balls in the urn is $\theta = 0.6$. The probability of drawing one white ball in a sample of one is then $p(W; \theta = 0.6) = 0.6$. We can also compute the probability for other outcomes, such as drawing two white balls in a sample of two, using the rules of the probability calculus (see section 1 of Hájek 2011). Generally, the probability that n_w white balls are drawn in a sample of size N , is given by the binomial distribution:

$$p(n_w; \theta = x) = \binom{N}{n_w} x^{n_w} (1 - x)^{(N - n_w)}$$

This is a specific example of a “sampling distribution”, $p(E | H)$, which gives the probability of certain evidence E in a sample, on the assumption that a certain hypothesis H is true. Calculation of the sampling distribution can in general be done *a priori*, given the rules of the probability calculus.

However, the problem of induction is the inverse problem. We want to infer not what the sample will be like, with a known hypothesis, rather we want to infer a hypothesis about the general situation or population, based on the observation of a limited sample. The probabilities of the candidate hypotheses can then be used to inform predictions about further observations. In the case of the urn, for example, we want to know what the observation of a particular sample frequency of white balls, $\frac{n_w}{N}$, tells us about θ , the proportion of white balls in the urn.

The idea of the Bayesian approach is to assign probabilities not only to the events which constitute evidence, but also to hypotheses. One starts with a “prior probability” distribution over the relevant hypotheses $p(H)$. On learning some evidence E , the Bayesian updates the prior $p(H)$ to the conditional probability $p(H | E)$. This update rule is called the “rule of conditionalisation”. The conditional probability $p(H | E)$ is known as the “posterior probability”, and is calculated using Bayes’ rule:

$$p(H | E) = \frac{p(E | H)p(H)}{p(E)}$$

Here the sampling distribution can be taken to be a conditional probability $p(E | H)$, which is known as the “likelihood” of the hypothesis H on evidence E .

One can then go on to compute the predictive distribution for as yet unobserved data E' , given observations E . The predictive distribution in a Bayesian approach is given by

$$p(E' | E) = \sum_H p(E' | H)p(H | E)$$

where the sum becomes an integral in cases where H is a continuous variable.

For the urn example, we can compute the posterior probability $p(\theta | n_w)$ using Bayes’ rule, and the likelihood given by the binomial distribution above. In order to do so, we also need to assign a prior probability distribution to the parameter θ . One natural choice, which was made early on by Bayes himself and by Laplace, is to put a uniform prior over the parameter θ . Bayes’ own rationale for this choice was that then if you work out the probability of each value for the number of whites in the sample based only on the prior, before any data is observed, all those probabilities are equal. Laplace had a different justification, based on the Principle of Indifference. This principle states that if you don’t have any reason to favor one hypothesis over another, you should assign them all equal probabilities.

With the choice of uniform prior, the posterior probability and predictive distribution can be calculated. It turns out that the probability that the next ball will be white, given that n_w of N draws were white, is given by

$$p(w | n_w) = \frac{n_w + 1}{N + 2}$$

This is Laplace’s famous “rule of succession” (1814). Suppose on the basis of observing 90 white balls out of 100, we calculate by the rule of succession that the probability of the next ball being white is $91/102 = 0.89$. It is quite conceivable that the next ball might be black. Even in the case, where all 100 balls have been white, so that the probability of the next ball being white is 0.99, there is still a small probability that the next ball is not white. What the probabilistic reasoning supplies then is not an argument to the conclusion that the next ball will be a certain color, but an argument to the conclusion that certain future observations are very *likely* given what has been observed in the past.

Overall, the Bayes-Laplace argument in the urn case provides an example of how probabilistic reasoning can take us from evidence about observations in the past to a prediction for how likely certain future observations are. The question is what kind of solution, if any, this type of calculation provides to the problem of induction. At first sight, since it is just a mathematical calculation, it looks as though it does indeed provide an *a priori* argument from the premises of an inductive inference to the proposition that a certain conclusion is probable.

However, in order to establish this definitively, one would need to argue that all the components and assumptions of the argument are *a priori* and this requires further examination of at least three important issues.

First, the Bayes-Laplace argument relies on the rules of the probability calculus. What is the status of these rules? Does following them amount to *a priori* reasoning? The answer to this depends in part on how probability itself is interpreted. Broadly speaking, there are prominent interpretations of probability according to which the rules plausibly have *a priori* status and could form the basis of a demonstrative argument. These include the classical interpretation originally developed by Laplace (1814), the logical

interpretation (Keynes (1921), Johnson (1921), Jeffreys (1939), Carnap (1950), Cox (1946, 1961), and the subjectivist interpretation of Ramsey (1926), Savage (1954), and de Finetti (1964). Attempts to argue for a probabilistic *a priori* solution to the problem of induction have been primarily associated with these interpretations.

Secondly, in the case of the urn, the Bayes-Laplace argument is based on a particular probabilistic model—the binomial model. This involves the assumption that there is a parameter describing an unknown proportion θ of balls in the urn, and that the data amounts to independent draws from a distribution over that parameter. What is the basis of these assumptions? Do they generalize to other cases beyond the actual urn case—i.e., can we see observations in general as analogous to draws from an “Urn of Nature”? There has been a persistent worry that these types of assumptions, while reasonable when applied to the case of drawing balls from an urn, will not hold for other cases of inductive inference. Thus, the probabilistic solution to the problem of induction might be of relatively limited scope. At the least, there are some assumptions going into the choice of model here that need to be made explicit. Arguably the choice of model introduces empirical assumptions, which would mean that the probabilistic solution is not an *a priori* one.

Thirdly, the Bayes-Laplace argument relies on a particular choice of prior probability distribution. What is the status of this assignment, and can it be based on *a priori* principles? Historically, the Bayes-Laplace choice of a uniform prior, as well as the whole concept of classical probability, relied on the Principle of Indifference. This principle has been regarded by many as an *a priori* principle. However, it has also been subjected to much criticism on the grounds that it can give rise to inconsistent probability assignments (Bertrand 1888; Borel 1909; Keynes 1921). Such inconsistencies are produced by there being more than one way to carve up the space of alternatives, and different choices give rise to conflicting probability assignments. One attempt to rescue the Principle of Indifference has been to appeal to explanationism, and argue that the principle should be applied only to the carving of the space at “the most explanatorily basic level”, where this level is identified according to an *a priori* notion of explanatory priority (Huemer 2009).

The quest for an *a priori* argument for the assignment of the prior has been largely abandoned. For many, the subjectivist foundations developed by Ramsey, de Finetti and Savage provide a more satisfactory basis for understanding probability. From this point of view, it is a mistake to try to introduce any further *a priori* constraints on the probabilities beyond those dictated by the probability rules themselves. Rather the assignment of priors may reflect personal opinions or background knowledge, and no prior is *a priori* an unreasonable choice.

So far, we have considered probabilistic arguments which place probabilities over hypotheses in a hypothesis space as well as observations. There is also a tradition of attempts to determine what probability distributions we should have, given certain observations, from the starting point of a joint probability distribution over all the observable variables. One may then postulate axioms directly on this distribution over observables, and examine the consequences for the predictive distribution. Much of the development of inductive logic, including the influential programme by Carnap, proceeded in this manner (Carnap 1950, 1952).

This approach helps to clarify the role of the assumptions behind probabilistic models. One assumption that one can make about the observations is that they are “exchangeable”. This means that the joint distribution of the random variables is invariant under permutations. Informally, this means that the order of the observations does not affect the probability. For instance, in the urn case, this would mean that drawing first a white ball and then a black ball is just as probable as first drawing a black and then a white. De Finetti proved a general representation theorem that if the joint probability distribution of an

infinite sequence of random variables is assumed to be exchangeable, then it can be written as a mixture of distribution functions from each of which the data behave as if they are independent random draws (de Finetti 1964). In the case of the urn example, the theorem shows that it is *as if* the data are independent random draws from a binomial distribution over a parameter θ , which itself has a prior probability distribution.

The assumption of exchangeability may be seen as a natural formalization of Hume's assumption that the past resembles the future. This is intuitive because assuming exchangeability means thinking that the order of observations, both past and future, does not matter to the probability assignments.

However, the development of the programme of inductive logic revealed that many generalizations are possible. For example, Johnson proposed to assume an axiom he called the "sufficientness postulate". This states that outcomes can be of a number of different types, and that the conditional probability that the next outcome is of type i depends only on the number of previous trials and the number of previous outcomes of type i (Johnson 1932). Assuming the sufficientness postulate for three or more types gives rise to a general predictive distribution corresponding to Carnap's "continuum of inductive methods" (Carnap 1952). This predictive distribution takes the form:

$$p(i \mid N_1, N_2, \dots, N_t) = \frac{N_i + k}{N_1 + N_2 + \dots + N_t + kt}$$

for some positive number k . This reduces to Laplace's rule of succession when $t = 2$ and $k = 1$.

Generalizations of the notion of exchangeability, such as "partial exchangeability" and "Markov exchangeability", have been explored, and these may be thought of as forms of symmetry assumption (Zabell 1988; Skyrms 2012). As less restrictive axioms on the probabilities for observables are assumed, the result is that there is no longer a unique result for the probability of a prediction, but rather a whole class of possible probabilities, mapped out by a generalized rule of succession such as the above. Therefore, in this tradition, as in the Bayes-Laplace approach, we have moved away from producing an argument which produces a unique a priori probabilistic answer to Hume's problem.

One might think then that the assignment of the prior, or the relevant corresponding postulates on the observable probability distribution, is precisely where empirical assumptions enter into inductive inferences. The probabilistic calculations are empirical arguments, rather than *a priori* ones. If this is correct, then the probabilistic framework has not in the end provided an *a priori* solution to the problem of induction, but it has rather allowed us to clarify what could be meant by Hume's claim that inductive inferences rely on the Uniformity Principle.

3.4 Partial solutions

Some think that although the problem of induction is not solved, there is in some sense a partial solution, which has been called a "logical solution". Howson, for example, argues that "*Inductive reasoning is justified to the extent that it is sound, given appropriate premises*" (Howson 2000: 239, his emphasis). According to this view, there is no getting away from an empirical premise for inductive inferences, but we might still think of Bayesian conditioning as functioning like a kind of logic or "consistency constraint" which "generates predictions from the assumptions and observations together" (Romeijn 2004: 360). Once we have an empirical assumption, instantiated in the prior probability, and the observations, Bayesian conditioning tells us what the resulting predictive probability distribution should be.

The idea of a partial solution also arises in the context of the learning theory that grounds contemporary machine learning. Machine learning is a field in computer science concerned with algorithms that learn from experience. Examples are algorithms which can be trained to recognise or classify patterns in data. Learning theory concerns itself with finding mathematical theorems which guarantee the performance of algorithms which are in practical use. In this domain, there is a well-known finding that learning algorithms are only effective if they have ‘inductive bias’ — that is, if they make some *a priori* assumptions about the domain they are employed upon (Mitchell 1997).

The idea is also given formal expression in the so-called ‘No-Free-Lunch theorems’ (Wolpert 1992, 1996, 1997). These can be interpreted as versions of the argument in Hume’s first fork since they establish that there can be no contradiction in the algorithm not performing well, since there are *a priori* possible situations in which it does not (Sterkenburg and Grünwald 2021:9992). Given Hume’s premise [P3](#), this rules out a demonstrative argument for its good performance.

Premise [P3](#) can perhaps be challenged on the grounds that *a priori* justifications can also be given for contingent propositions. Even though an inductive inference can fail in some possible situations, it could still be reasonable to form an expectation of reliability if we spread our credence equally over all the possibilities and have reason to think (or at least no reason to doubt) that the cases where inductive inference is unreliable require a ‘very specific arrangement of things’ and thus form a small fraction of the total space of possibilities (White 2015). The No-Free-Lunch theorems make difficulties for this approach since they show that if we put a uniform distribution over all logically possible sequences of future events, any learning algorithm is expected to have a generalisation error of 1/2, and hence to do no better than guessing at random (Schurz 2021b).

The No-Free-Lunch theorems may be seen as fundamental limitations on justifying learning algorithms when these algorithms are seen as ‘purely data-driven’ — that is as mappings from possible data to conclusions. However, learning algorithms may also be conceived as functions not only of input data, but also of a particular model (Sterkenburg and Grünwald 2021). For example, the Bayesian ‘algorithm’ gives a universal recipe for taking a particular model and prior and updating on the data. A number of theorems in learning theory provide general guarantees for the performance of such recipes. For instance, there are theorems which guarantee convergence of the Bayesian algorithm (Ghosal, Ghosh and van der Vaart 2000, Ghosal, Lember and van der Vaart 2008). In each instantiation, this convergence is relative to a particular specific prior. Thus, although the considerations first raised by Hume, and later instantiated in the No-Free-Lunch theorems, preclude any universal model-independent justification for learning algorithms, it does not rule out partial justifications in the form of such general *a priori* ‘model-relative’ learning guarantees (Sterkenburg and Grünwald 2021).

3.5 The combinatorial approach

An alternative attempt to use probabilistic reasoning to produce an *a priori* justification for inductive inferences is the so-called “combinatorial” solution. This was first put forward by Donald C. Williams (1947) and later developed by David Stove (1986).

Like the Bayes-Laplace argument, the solution relies heavily on the idea that straightforward *a priori* calculations can be done in a “direct inference” from population to sample. As we have seen, given a certain population frequency, the probability of getting different frequencies in a sample can be calculated straightforwardly based on the rules of the probability calculus. The Bayes-Laplace argument relied on inverting the probability distribution using Bayes’ rule to get from the sampling distribution to the posterior distribution. Williams instead proposes that the inverse inference may be based on a certain logical syllogism: the proportional (or statistical) syllogism.

The proportional, or statistical syllogism, is the following:

1. Of all the things that are M , m/n are P .
2. a is an M

Therefore, a is P , with probability m/n .

For example, if 90% of rabbits in a population are white and we observe a rabbit a , then the proportional syllogism says that we infer that a is white with a probability of 90%. Williams argues that the proportional syllogism is a non-deductive logical syllogism, which effectively interpolates between the syllogism for entailment

1. All M s are P
2. a is an M

Therefore, a is P .

And the syllogism for contradiction

1. No M is P
2. a is M

Therefore, a is not P .

This syllogism can be combined with an observation about the behavior of increasingly large samples. From calculations of the sampling distribution, it can be shown that as the sample size increases, the probability that the sample frequency is in a range which closely approximates the population frequency also increases. In fact, Bernoulli's law of large numbers states that the probability that the sample frequency approximates the population frequency tends to one as the sample size goes to infinity. Williams argues that such results support a "general over-all premise, common to all inductions, that samples 'match' their populations" (Williams 1947: 78).

We can then apply the proportional syllogism to samples from a population, to get the following argument:

1. Most samples match their population
2. S is a sample.

Therefore, S matches its population, with high probability.

This is an instance of the proportional syllogism, and it uses the general result about samples matching populations as the first major premise.

The next step is to argue that if we observe that the sample contains a proportion of m/n F s, then we can conclude that since this sample with high probability matches its population, the population, with high probability, has a population frequency that approximates the sample frequency m/n . Both Williams and Stove claim that this amounts to a logical *a priori* solution to the problem of induction.

A number of authors have expressed the view that the Williams-Stove argument is only valid if the sample S is drawn randomly from the population of possible samples—i.e., that any sample is as likely to be drawn as any other (Brown 1987; Will 1948; Giaquinto 1987). Sometimes this is presented as an objection to the application of the proportional syllogism. The claim is that the proportional syllogism is

only valid if a is drawn randomly from the population of M s. However, the response has been that there is no need to know that the sample is randomly drawn in order to apply the syllogism (Maher 1996; Campbell 2001; Campbell & Franklin 2004). Certainly if you have reason to think that your sampling procedure is more likely to draw certain individuals than others—for example, if you know that you are in a certain location where there are more of a certain type—then you should not apply the proportional syllogism. But if you have no such reasons, the defenders claim, it is quite rational to apply it. Certainly it is always possible that you draw an unrepresentative sample—meaning one of the few samples in which the sample frequency does not match the population frequency—but this is why the conclusion is only probable and not certain.

The more problematic step in the argument is the final step, which takes us from the claim that samples match their populations with high probability to the claim that having seen a particular sample frequency, the population from which the sample is drawn has frequency close to the sample frequency with high probability. The problem here is a subtle shift in what is meant by “high probability”, which has formed the basis of a common misreading of Bernoulli’s theorem. Hacking (1975: 156–59) puts the point in the following terms. Bernoulli’s theorem licenses the claim that much more often than not, a small interval around the sample frequency will include the true population frequency. In other words, it is highly probable in the sense of “usually right” to say that the sample matches its population. But this does not imply that the proposition that a small interval around the sample will contain the true population frequency is highly probable in the sense of “credible on each occasion of use”. This would mean that for any given sample, it is highly credible that the sample matches its population. It is quite compatible with the claim that it is “usually right” that the sample matches its population to say that there are some samples which do not match their populations at all. Thus one cannot conclude from Bernoulli’s theorem that for any given sample frequency, we should assign high probability to the proposition that a small interval around the sample frequency will contain the true population frequency. But this is exactly the slide that Williams makes in the final step of his argument. Maher (1996) argues in a similar fashion that the last step of the Williams-Stove argument is fallacious. In fact, if one wants to draw conclusions about the probability of the population frequency given the sample frequency, the proper way to do so is by using the Bayesian method described in the previous section. But, as we there saw, this requires the assignment of prior probabilities, and this explains why many people have thought that the combinatorial solution somehow illicitly presupposed an assumption like the principle of indifference. The Williams-Stove argument does not in fact give us an alternative way of inverting the probabilities which somehow bypasses all the issues that Bayesians have faced.

4. Tackling the Second Horn of Hume’s Dilemma

So far we have considered ways in which the first horn of Hume’s dilemma might be tackled. But it is of course also possible to take on the second horn instead.

One may argue that a probable argument would not, despite what Hume says, be circular in a problematic way (we consider responses of this kind in [section 4.1](#)). Or, one might attempt to argue that probable arguments are not circular at all ([section 4.2](#)).

4.1 Inductive Justifications of Induction

One way to tackle the second horn of Hume’s dilemma is to reject [premise P6](#), which rules out circular arguments. Some have argued that certain kinds of circular arguments would provide an acceptable justification for the inductive inference. Since the justification would then itself be an inductive one, this approach is often referred to as an “inductive justification of induction”.

First we should examine how exactly the Humean circularity supposedly arises. Take the simple case of enumerative inductive inference that follows the following pattern (X):

Most observed F s have been G s

Therefore: Most F s are G s.

Hume claims that such arguments presuppose the Uniformity Principle (UP). According to premises [P7](#) and [P8](#), this supposition also needs to be supported by an argument in order that the inductive inference be justified. A natural idea is that we can argue for the Uniformity Principle on the grounds that “it works”. We know that it works, because past instances of arguments which relied upon it were found to be successful. This alone however is not sufficient unless we have reason to think that such arguments will also be successful in the future. That claim must itself be supported by an inductive argument (S):

Most arguments of form X that rely on UP have succeeded in the past.

Therefore, most arguments of form X that rely on UP succeed.

But this argument itself depends on the UP, which is the very supposition which we were trying to justify.

As we have seen in [section 2](#), some reject Hume’s claim that all inductive inferences presuppose the UP. However, the argument that basing the justification of the inductive inference on a probable argument would result in circularity need not rely on this claim. The circularity concern can be framed more generally. If argument S relies on *something* which is already presupposed in inference X , then argument S cannot be used to justify inference X . The question though is what precisely the something is.

Some authors have argued that in fact S does not rely on any premise or even presupposition that would require us to already know the conclusion of X . S is then not a “premise circular” argument. Rather, they claim, it is “rule-circular”—it relies on a rule of inference in order to reach the conclusion that that very rule is reliable. Suppose we adopt the rule R which says that when it is observed that most F s are G s, we should infer that most F s are G s. Then inference X relies on rule R . We want to show that rule R is reliable. We could appeal to the fact that R worked in the past, and so, by an inductive argument, it will also work in the future. Call this argument S^* :

Most inferences following rule R have been successful

Therefore, most inferences following R are successful.

Since this argument itself uses rule R , using it to establish that R is reliable is rule-circular.

Some authors have then argued that although premise-circularity is vicious, rule-circularity is not (Cleve 1984; Papineau 1992). One reason for thinking rule-circularity is not vicious would be if it is not necessary to know or even justifiably believe that rule R is reliable in order to move to a justified conclusion using the rule. This is a claim made by externalists about justification (Cleve 1984). They say that as long as R is *in fact* reliable, one can form a justified belief in the conclusion of an argument relying on R , as long as one has justified belief in the premises.

If one is not persuaded by the externalist claim, one might attempt to argue that rule circularity is benign in a different fashion. For example, the requirement that a rule be shown to be reliable without any rule-circularity might appear unreasonable when the rule is of a very fundamental nature. As Lange puts it:

It might be suggested that although a circular argument is ordinarily unable to justify its conclusion, a circular argument is acceptable in the case of justifying a fundamental form of reasoning. After all,

there is nowhere more basic to turn, so all that we can reasonably demand of a fundamental form of reasoning is that it endorse itself. (Lange 2011: 56)

Proponents of this point of view point out that even deductive inference cannot be justified deductively. Consider Lewis Carroll's dialogue between Achilles and the Tortoise (Carroll 1895). Achilles is arguing with a Tortoise who refuses to perform *modus ponens*. The Tortoise accepts the premise that p , and the premise that p implies q but he will not accept q . How can Achilles convince him? He manages to persuade him to accept another premise, namely "if p and p implies q , then q ". But the Tortoise is still not prepared to infer to q . Achilles goes on adding more premises of the same kind, but to no avail. It appears then that *modus ponens* cannot be justified to someone who is not already prepared to use that rule.

It might seem odd if premise circularity were vicious, and rule circularity were not, given that there appears to be an easy interchange between rules and premises. After all, a rule can always, as in the Lewis Carroll story, be added as a premise to the argument. But what the Carroll story also appears to indicate is that there is indeed a fundamental difference between being prepared to accept a premise stating a rule (the Tortoise is happy to do this), and being prepared to use that rule (this is what the Tortoise refuses to do).

Suppose that we grant that an inductive argument such as S (or S^*) can support an inductive inference X without vicious circularity. Still, a possible objection is that the argument simply does not provide a full justification of X . After all, less sane inference rules such as counterinduction can support themselves in a similar fashion. The counterinductive rule is CI:

Most observed As are Bs .

Therefore, it is not the case that most As are Bs .

Consider then the following argument CI*:

Most CI arguments have been unsuccessful

Therefore, it is not the case that most CI arguments are unsuccessful, i.e., many CI arguments are successful.

This argument therefore establishes the reliability of CI in a rule-circular fashion (see Salmon 1963).

Argument S can be used to support inference X , but only for someone who is already prepared to infer inductively by using S . It cannot convince a skeptic who is not prepared to rely upon that rule in the first place. One might think then that the argument is simply not achieving very much.

The response to these concerns is that, as Papineau puts it, the argument is "not *supposed* to do very much" (Papineau 1992: 18). The fact that a counterinductivist counterpart of the argument exists is true, but irrelevant. It is conceded that the argument cannot persuade either a counterinductivist, or a skeptic. Nonetheless, proponents of the inductive justification maintain that there is still some added value in showing that inductive inferences are reliable, even when we already accept that there is nothing problematic about them. The inductive justification of induction provides a kind of important consistency check on our existing beliefs.

4.2 No Rules

It is possible to go even further in an attempt to dismantle the Humean circularity. Maybe inductive inferences do not even have a rule in common. What if every inductive inference is essentially unique? This can be seen as rejecting Hume's premise [P5](#). Okasha, for example, argues that Hume's circularity

problem can be evaded if there are “no rules” behind induction (Okasha 2005a,b). Norton puts forward the similar idea that all inductive inferences are material, and have nothing formal in common (Norton 2003, 2010, 2021).

Proponents of such views have attacked Hume’s claim that there is a UP on which all inductive inferences are based. There have long been complaints about the vagueness of the Uniformity Principle (Salmon 1953). The future only resembles the past in some respects, but not others. Suppose that on all my birthdays so far, I have been under 40 years old. This does not give me a reason to expect that I will be under 40 years old on my next birthday. There seems then to be a major lacuna in Hume’s account. He might have explained or described how we draw an inductive inference, on the assumption that it is one we *can* draw. But he leaves untouched the question of how we distinguish between cases where we extrapolate a regularity legitimately, regarding it as a law, and cases where we do not.

Nelson Goodman is often seen as having made this point in a particularly vivid form with his “new riddle of induction” (Goodman 1955: 59–83). Suppose we define a predicate “grue” in the following way. An object is “grue” when it is green if observed before time *t* and blue otherwise. Goodman considers a thought experiment in which we observe a bunch of green emeralds before time *t*. We could describe our results by saying all the observed emeralds are green. Using a simple enumerative inductive schema, we could infer from the result that all observed emeralds are green, that all emeralds are green. But equally, we could describe the same results by saying that all observed emeralds are grue. Then using the same schema, we could infer from the result that all observed emeralds are grue, that all emeralds are grue. In the first case, we expect an emerald observed after time *t* to be green, whereas in the second, we expect it to be blue. Thus the two predictions are incompatible. Goodman claims that what Hume omitted to do was to give any explanation for why we project predicates like “green”, but not predicates like “grue”. This is the “new riddle”, which is often taken to be a further problem of induction that Hume did not address.

One moral that could be taken from Goodman is that there is not one general Uniformity Principle that all probable arguments rely upon (Sober 1988; Norton 2003; Okasha 2001, 2005a,b, Jackson 2019). Rather each inductive inference presupposes some more specific empirical presupposition. A particular inductive inference depends on some specific way in which the future resembles the past. It can then be justified by another inductive inference which depends on some quite different empirical claim. This will in turn need to be justified—by yet another inductive inference. The nature of Hume’s problem in the second horn is thus transformed. There is no circularity. Rather there is a regress of inductive justifications, each relying on their own empirical presuppositions (Sober 1988; Norton 2003; Okasha 2001, 2005a,b).

One way to put this point is to say that Hume’s argument rests on a quantifier shift fallacy (Sober 1988; Okasha 2005a). Hume says that there exists a general presupposition for all inductive inferences, whereas he should have said that for each inductive inference, there is some presupposition. Different inductive inferences then rest on different empirical presuppositions, and the problem of circularity is evaded.

What will then be the consequence of supposing that Hume’s problem should indeed have been a regress, rather than a circularity? Here different opinions are possible. On the one hand, one might think that a regress still leads to a skeptical conclusion (Schurz and Thorn 2020). So although the exact form in which Hume stated his problem was not correct, the conclusion is not substantially different (Sober 1988). Another possibility is that the transformation mitigates or even removes the skeptical problem. For example, Norton argues that the upshot is a dissolution of the problem of induction, since the regress of justifications benignly terminates (Norton 2003). And Okasha more mildly suggests that even if the regress is infinite, “Perhaps infinite regresses are less bad than vicious circles after all” (Okasha 2005b: 253).

Any dissolution of Hume’s circularity does not depend only on arguing that the UP should be replaced by empirical presuppositions which are specific to each inductive inference. It is also necessary to establish that inductive inferences share no common rules—otherwise there will still be at least some rule-circularity. Okasha suggests that the Bayesian model of belief-updating is an illustration how induction can be characterized in a rule-free way, but this is problematic, since in this model all inductive inferences still share the common rule of Bayesian conditionalisation. Norton’s material theory of induction postulates a rule-free characterization of induction, but it is not clear whether it really can avoid any role for general rules (Achinstein 2010, Kelly 2010, Worrall 2010).

5. Alternative Conceptions of Justification

Hume is usually read as delivering a negative verdict on the possibility of justifying inference *I*, via a premise such as [P8](#), though as we have seen in [section 2](#), some have questioned whether Hume is best interpreted as drawing a conclusion about justification of inference *I* at all. In this section we examine approaches which question in different ways whether [premise P8](#) really does give a valid necessary condition for justification of inference *I* and propose various alternative conceptions of justification.

5.1 Postulates and Hinges

One approach has been to turn to general reflection on what is even needed for justification of an inference in the first place. For example, Wittgenstein raised doubts over whether it is even meaningful to ask for the grounds for inductive inferences.

If anyone said that information about the past could not convince him that something would happen in the future, I should not understand him. One might ask him: what do you expect to be told, then? What sort of information do you call a ground for such a belief? ... If these are not grounds, then what are grounds?—If you say these are not grounds, then you must surely be able to state what must be the case for us to have the right to say that there are grounds for our assumption....
(Wittgenstein 1953: 481)

One might not, for instance, think that there even needs to be a chain of reasoning in which each step or presupposition is supported by an argument. Wittgenstein took it that there are some principles so fundamental that they do not require support from any further argument. They are the “hinges” on which enquiry turns.

Out of Wittgenstein’s ideas has developed a general notion of “entitlement”, which is a kind of rational warrant to hold certain propositions which does not come with the same requirements as “justification”. Entitlement provides epistemic rights to hold a proposition, without responsibilities to base the belief in it on an argument. Crispin Wright (2004) has argued that there are certain principles, including the Uniformity Principle, that we are entitled in this sense to hold.

Some philosophers have set themselves the task of determining a set or sets of postulates which form a plausible basis for inductive inferences. Bertrand Russell, for example, argued that five postulates lay at the root of inductive reasoning (Russell 1948). Arthur Burks, on the other hand, proposed that the set of postulates is not unique, but there may be multiple sets of postulates corresponding to different inductive methods (Burks 1953, 1955).

The main objection to all these views is that they do not really solve the problem of induction in a way that adequately secures the pillars on which inductive inference stands. As Salmon puts it, “admission of

unjustified and unjustifiable postulates to deal with the problem is tantamount to making scientific method a matter of faith” (Salmon 1966: 48).

5.2 Ordinary Language Dissolution

Rather than allowing undefended empirical postulates to give normative support to an inductive inference, one could instead argue for a completely different conception of what is involved in justification. Like Wittgenstein, later ordinary language philosophers, notably P.F. Strawson, also questioned what exactly it means to ask for a justification of inductive inferences (Strawson 1952). This has become known as the “Ordinary language dissolution” of the problem of induction.

Strawson points out that it could be meaningful to ask for a deductive justification of inductive inferences. But it is not clear that this is helpful since this is effectively “a demand that induction shall be shown to be really a kind of deduction” (Strawson 1952: 230). Rather, Strawson says, when we ask about whether a particular inductive inference is justified, we are typically judging whether it conforms to our usual inductive standards. Suppose, he says, someone has formed the belief by inductive inference that All f 's are g . Strawson says that if that person is asked for their grounds or reasons for holding that belief,

I think it would be felt to be a satisfactory answer if he replied: “Well, in all my wide and varied experience I’ve come across innumerable cases of f and never a case of f which wasn’t a case of g ”. In saying this, he is clearly claiming to have *inductive* support, *inductive* evidence, of a certain kind, for his belief. (Strawson 1952)

That is just because inductive support, as it is usually understood, simply consists of having observed many positive instances in a wide variety of conditions.

In effect, this approach denies that producing a chain of reasoning is a necessary condition for justification. Rather, an inductive inference is justified if it conforms to the usual standards of inductive justification. But, is there more to it? Might we not ask what reason we have to rely on those inductive standards?

It surely makes sense to ask whether a particular inductive inference is justified. But the answer to that is fairly straightforward. Sometimes people have enough evidence for their conclusions and sometimes they do not. Does it also make sense to ask about whether inductive procedures generally are justified? Strawson draws the analogy between asking whether a particular act is legal. We may answer such a question, he says, by referring to the law of the land.

But it makes no sense to inquire in general whether the law of the land, the legal system as a whole, is or is not legal. For to what legal standards are we appealing? (Strawson 1952: 257)

According to Strawson,

It is an analytic proposition that it is reasonable to have a degree of belief in a statement which is proportional to the strength of the evidence in its favour; and it is an analytic proposition, though not a proposition of mathematics, that, other things being equal, the evidence for a generalisation is strong in proportion as the number of favourable instances, and the variety of circumstances in which they have been found, is great. So to ask whether it is reasonable to place reliance on inductive procedures is like asking whether it is reasonable to proportion the degree of one’s convictions to the strength of the evidence. Doing this is what “being reasonable” *means* in such a context. (Strawson 1952: 256–57)

Thus, according to this point of view, there is no further question to ask about whether it is reasonable to rely on inductive inferences.

The ordinary language philosophers do not explicitly argue against Hume's [premise P8](#). But effectively what they are doing is offering a whole different story about what it would mean to be justified in believing the conclusion of inductive inferences. What is needed is just conformity to inductive standards, and there is no real meaning to asking for any further justification for those.

The main objection to this view is that conformity to the usual standards is insufficient to provide the needed justification. What we need to know is whether belief in the conclusion of an inductive inference is “epistemically reasonable or justified in the sense that ...there is reason to think that it is likely to be true” (BonJour 1998: 198). The problem Hume has raised is whether, despite the fact that inductive inferences have tended to produce true conclusions in the past, we have reason to think the conclusion of an inductive inference we now make is likely to be true. Arguably, establishing that an inductive inference is rational in the sense that it follows inductive standards is not sufficient to establish that its conclusion is likely to be true. In fact Strawson allows that there is a question about whether “induction will continue to be successful”, which is distinct from the question of whether induction is rational. This question he does take to hinge on a “contingent, factual matter” (Strawson 1952: 262). But if it is this question that concerned Hume, it is no answer to establish that induction is rational, unless that claim is understood to involve or imply that an inductive inference carried out according to rational standards is likely to have a true conclusion.

5.3 Pragmatic vindication of induction

Another solution based on an alternative criterion for justification is the “pragmatic” approach initiated by Reichenbach (1938 [2006]). Reichenbach did think Hume's argument unassailable, but nonetheless he attempted to provide a weaker kind of justification for induction. In order to emphasize the difference from the kind of justification Hume sought, some have given it a different term and refer to Reichenbach's solution as a “vindication”, rather than a justification of induction (Feigl 1950; Salmon 1963).

Reichenbach argued that it was not necessary for the justification of inductive inference to show that its conclusion is true. Rather “the proof of the truth of the conclusion is only a sufficient condition for the justification of induction, not a necessary condition” (Reichenbach 2006: 348). If it could be shown, he says, that inductive inference is a necessary condition of success, then even if we do not know that it will succeed, we still have some reason to follow it. Reichenbach makes a comparison to the situation where a man is suffering from a disease, and the physician says “I do not know whether an operation will save the man, but if there is any remedy, it is an operation” (Reichenbach 1938 [2006: 349]). This provides some kind of justification for operating on the man, even if one does not know that the operation will succeed.

In order to get a full account, of course, we need to say more about what is meant for a method to have “success”, or to “work”. Reichenbach thought that this should be defined in relation to the aim of induction. This aim, he thought, is “*to find series of events whose frequency of occurrence converges towards a limit*” (1938 [2006: 350]).

Reichenbach applied his strategy to a general form of “statistical induction” in which we observe the relative frequency f_n of a particular event in n observations and then form expectations about the frequency that will arise when more observations are made. The “inductive principle” then states that if after a certain number of instances, an observed frequency of m/n is observed, for any prolongation of the series of observations, the frequency will continue to fall within a small interval of m/n . Hume's

examples are special cases of this principle, where the observed frequency is 1. For example, in Hume's bread case, suppose bread was observed to nourish n times out of n (i.e. an observed frequency of 100%), then according to the principle of induction, we expect that as we observe more instances, the frequency of nourishing ones will continue to be within a very small interval of 100%. Following this inductive principle is also sometimes referred to as following the "straight rule". The problem then is to justify the use of this rule.

Reichenbach argued that even if Hume is right to think that we cannot be justified in thinking for any particular application of the rule that the conclusion is likely to be true, for the purposes of practical action we do not need to establish this. We can instead regard the inductive rule as resulting in a "posit", or statement that we deal with as if it is true. We posit a certain frequency f on the basis of our evidence, and this is like making a wager or bet that the frequency is in fact f . One strategy for positing frequencies is to follow the rule of induction.

Reichenbach proposes that we can show that the rule of induction meets his weaker justification condition. This does not require showing that following the inductive principle will always work. It is possible that the world is so disorderly that we cannot construct series with any limits. In that case, neither the inductive principle, nor any other method will succeed. But, he argues, if there is a limit, by following the inductive principle we will eventually find it. There is some element of a series of observations, beyond which the principle of induction will lead to the true value of the limit. Although the inductive rule may give quite wrong results early in the sequence, as it follows chance fluctuations in the sample frequency, it is guaranteed to eventually approximate the limiting frequency, if such a limit exists. Therefore, the rule of induction is justified as an instrument of positing because it is a method of which we know that if it is possible to achieve the aim of inductive inference we shall do so by means of this method (Reichenbach 1949: 475).

One might question whether Reichenbach has achieved his goal of showing that following the inductive rule is a necessary condition of success. In order to show that, one would also need to establish that no other methods can also achieve the aim. But, as Reichenbach himself recognises, many other rules of inference as well as the straight rule may also converge on the limit (Salmon 1966: 53). In fact, any method which converges asymptotically to the straight rule also does so. An easily specified class of such rules are those which add to the inductive rule a function c_n in which the c_n converge to zero with increasing n .

Reichenbach makes two suggestions aimed at avoiding this problem. On the one hand, he claims, since we have no real way to pick between methods, we might as well just use the inductive rule since it is "easier to handle, owing to its descriptive simplicity". He also claims that the method which embodies the "smallest risk" is following the inductive rule (Reichenbach 1938 [2006: 355–356]).

There is also the concern that there could be a completely different kind of rule which converges on the limit. We can consider, for example, the possibility of a soothsayer or psychic who is able to predict future events reliably. Here Reichenbach argues that induction is still necessary in such a case, because it has to be used to check whether the other method works. It is only by using induction, Reichenbach says, that we could recognise the reliability of the alternative method, by examining its track record.

In assessing this argument, it is helpful to distinguish between levels at which the principle of induction can be applied. Following Skyrms (2000), we may distinguish between level 1, where candidate methods are applied to ordinary events or individuals, and level 2, where they are applied not to individuals or events, but to the arguments on level 1. Let us refer to "object-induction" when the inductive principle is applied at level 1, and "meta-induction" when it is applied at level 2. Reichenbach's response does not rule out the possibility that another method might do better than object-induction at level 1. It only shows

that the success of that other method may be recognised by a meta-induction at level 2 (Skyrms 2000). Nonetheless, Reichenbach's thought was later picked up and developed into the suggestion that a meta-inductivist who applies induction not only at the object level to observations, but also to the success of others' methods, might by those means be able to do as well predictively as the alternative method (Schurz 2008; see [section 5.5](#) for more discussion of meta-induction).

Reichenbach's justification is generally taken to be a pragmatic one, since though it does not supply knowledge of a future event, it supplies a sufficient reason for action (Reichenbach 1949: 481). One might question whether a pragmatic argument can really deliver an all-purpose, general justification for following the inductive rule. Surely a pragmatic solution should be sensitive to differences in pay-offs that depend on the circumstances. For example, Reichenbach offers the following analogue to his pragmatic justification:

We may compare our situation to that of a man who wants to fish in an unexplored part of the sea. There is no one to tell him whether or not there are fish in this place. Shall he cast his net? Well, if he wants to fish in that place, I should advise him to cast the net, to take the chance at least. It is preferable to try even in uncertainty than not to try and be certain of getting nothing. (Reichenbach 1938 [2006: 362–363])

As Lange points out, the argument here “presumes that there is no cost to trying”. In such a situation, “the fisherman has everything to gain and nothing to lose by casting his net” (Lange 2011: 77). But if there is some significant cost to making the attempt, it may not be so clear that the most rational course of action is to cast the net. Similarly, whether or not it would make sense to adopt the policy of making no predictions, rather than the policy of following the inductive rule, may depend on what the practical penalties are for being wrong. A pragmatic solution may not be capable of offering rationale for following the inductive rule which is applicable in all circumstances.

Another question is whether Reichenbach has specified the aim of induction too narrowly. Finding series of events whose frequency of occurrence converges to a limit ties the vindication to the long-run, while allowing essentially no constraint on what can be posited in the short-run. Yet it is in the short run that inductive practice actually occurs and where it really needs justification (BonJour 1998: 194; Salmon 1966: 53).

5.4 Formal Learning Theory

Formal learning theory can be regarded as a kind of extension of the Reichenbachian programme. It does not offer justifications for inductive inferences in the sense of giving reasons why they should be taken as likely to provide a true conclusion. Rather it offers a “means-ends” epistemology -- it provides reasons for following particular methods based on their optimality in achieving certain desirable epistemic ends, even if there is no guarantee that at any given stage of inquiry the results they produce are at all close to the truth (Schulte 1999).

Formal learning theory is particularly concerned with showing that methods are “logically reliable” in the sense that they arrive at the truth given any sequence of data consistent with our background knowledge (Kelly 1996). However, it goes further than this. As we have just seen, one of the problems for Reichenbach was that there are too many rules which converge in the limit to the true frequency. Which one should we then choose in the short-run? Formal learning theory broadens Reichenbach's general strategy by considering what happens if we have other epistemic goals besides long-run convergence to the truth. In particular, formal learning theorists have considered the goal of getting to the truth as efficiently, or quickly, as possible, as well as the goal of minimising the number of mind-changes, or

retractions along the way. It has then been argued that the usual inductive method, which is characterised by a preference for simpler hypotheses (Occam's razor), can be justified since it is the unique method which meets the standards for getting to the truth in the long run as efficiently as possible, with a minimum number of retractions (Kelly 2007).

Steel (2010) has proposed that the Principle of Induction (understood as a rule which makes inductive generalisations along the lines of the Straight Rule) can be given a means-ends justification by showing that following it is both necessary and sufficient for logical reliability. The proof is an *a priori* mathematical one, thus it allegedly avoids the circularity of Hume's second horn. However, Steel also does not see the approach as an attempt to grasp Hume's first horn, since the proof is only relative to a certain choice of epistemic ends.

As with other results in formal learning theory, this solution is also only valid relative to a given hypothesis space and conception of possible sequences of data. For this reason, some have seen it as not addressing Hume's problem of giving grounds for a particular inductive inference (Howson 2011). An alternative attitude is that it does solve a significant part of Hume's problem (Steel 2010). There is a similar dispute over formal learning theory's treatment of Goodman's riddle (Chart 2000, Schulte 2017).

5.5 Meta-induction

Another approach to pursuing a broadly Reichenbachian programme is Gerhard Schurz's strategy based on meta-induction (Schurz 2008, 2017, 2019). Schurz draws a distinction between applying inductive methods at the level of events—so-called “object-level” induction (OI), and applying inductive methods at the level of competing prediction methods—so-called “meta-induction” (MI). Whereas object-level inductive methods make predictions based on the events which have been observed to occur, meta-inductive methods make predictions based on aggregating the predictions of different available prediction methods according to their success rates. Here, the success rate of a method is defined according to some precise way of scoring success in making predictions.

The starting point of the meta-inductive approach is that the aim of inductive inference is not just, as Reichenbach had it, finding long-run limiting frequencies, but also predicting successfully in both the long and short run. Even if Hume has precluded showing that the inductive method is reliable in achieving successful prediction, perhaps it can still be shown that it is “predictively optimal”. A method is “predictively optimal” if it succeeds best in making successful predictions out of all competing methods, no matter what data is received. Schurz brings to bear results from the regret-based learning framework in machine learning that show that there is a meta-inductive strategy that is predictively optimal among all predictive methods that are accessible to an epistemic agent (Cesa-Bianchi and Lugosi 2006, Schurz 2008, 2017, 2019). This meta-inductive strategy, which Schurz calls “wMI”, predicts a weighted average of the predictions of the accessible methods, where the weights are “attractivities”, which measure the difference between the method's own success rate and the success rate of wMI.

The main result is that the wMI strategy is long-run optimal in the sense that it converges to the maximum success rate of the accessible prediction methods. Worst-case bounds for short-run performance can also be derived. The optimality result forms the basis for an *a priori* means-ends justification for the use of wMI. Namely, the thought is, it is reasonable to use wMI, since it achieves the best success rates possible in the long run out of the given methods.

Schurz also claims that this *a priori* justification of wMI, together with the contingent fact that inductive methods have so far been much more successful than non-inductive methods, gives rise to an *a posteriori* non-circular justification of induction. Since wMI will achieve in the long run the maximal success rate of the available prediction methods, it is reasonable to use it. But as a matter of fact, object-inductive

prediction methods have been more successful than non-inductive methods so far. Therefore Schurz says “it is meta-inductively justified to favor object-inductivistic strategies in the future” (Schurz 2019: 85). This justification, he claims, is not circular because meta-induction has an *a priori* independent justification. The idea is that since it is *a priori* justified to use wMI, it is also *a priori* justified to use the maximally successful method at the object level. Since it turns out that that the maximally successful method is object-induction, then we have a non-circular *a posteriori* argument that it is reasonable to use object-induction.

Schurz’s original theorems on the optimality of wMI apply to the case where there are finitely many predictive methods. One point of discussion is whether this amounts to an important limitation on its claims to provide a full solution of the problem of induction. The question then is whether it is necessary that the optimality results be extended to an infinite, or perhaps an expanding pool of strategies (Eckhardt 2010, Sterkenburg 2019, Schurz 2021a).

Another important issue concerns what it means for object-induction to be “meta-inductively justified”. The meta-inductive strategy wMI and object-induction are clearly different strategies. They could result in different predictions tomorrow, if OI would stop working and another method would start to do better. In that case, wMI would begin to favour the other method, and wMI would start to come apart from OI. The optimality results provide a reason to follow wMI. How exactly does object-induction inherit that justification? At most, it seems that we get a justification for following OI on the next time-step, on the grounds that OI’s prediction approximately coincides with that of wMI (Sterkenburg 2020, Sterkenburg (forthcoming)). However, this requires a stronger empirical postulate than simply the observation that OI has been more successful than non-inductive methods. It also requires something like that “as a matter of empirical fact, the strategy OI has been so much more successful than its competitors, that the meta-inductivist attributes it such a large share of the total weight that its prediction (approximately) coincides with OI’s prediction” (Sterkenburg 2020: 538). Furthermore, even if we allow that the empirical evidence does back up such a strong claim, the issue remains that the meta-inductive justification is in support of following the strategy of meta-induction, not in support of the *strategy* of following OI (Sterkenburg (2020), sec. 3.3.2).

6. Living with Inductive Skepticism

So far we have considered the various ways in which we might attempt to solve the problem of induction by resisting one or other premise of Hume’s argument. Some philosophers have however seen his argument as unassailable, and have thus accepted that it does lead to inductive skepticism, the conclusion that inductive inferences cannot be justified. The challenge then is to find a way of living with such a radical-seeming conclusion. We appear to rely on inductive inference ubiquitously in daily life, and it is also generally thought that it is at the very foundation of the scientific method. Can we go on with all this, whilst still seriously thinking none of it is justified by any rational argument?

One option here is to argue, as does Nicholas Maxwell, that the problem of induction is posed in an overly restrictive context. Maxwell argues that the problem does not arise if we adopt a different conception of science than the ‘standard empiricist’ one, which he denotes ‘aim-oriented empiricism’ (Maxwell 2017).

Another option here is to think that the significance of the problem of induction is somehow restricted to a skeptical context. Hume himself seems to have thought along these lines. For instance he says:

Nature will always maintain her rights, and prevail in the end over any abstract reasoning whatsoever. Though we should conclude, for instance, as in the foregoing section, that, in all reasonings from experience, there is a step taken by the mind, which is not supported by any argument or process of the understanding; there is no danger, that these reasonings, on which almost all knowledge depends, will ever be affected by such a discovery. (E. 5.1.2)

Hume's purpose is clearly not to argue that we should not make inductive inferences in everyday life, and indeed his whole method and system of describing the mind in naturalistic terms depends on inductive inferences through and through. The problem of induction then must be seen as a problem that arises only at the level of philosophical reflection.

Another way to mitigate the force of inductive skepticism is to restrict its scope. Karl Popper, for instance, regarded the problem of induction as insurmountable, but he argued that science is not in fact based on inductive inferences at all (Popper 1935 [1959]). Rather he presented a deductivist view of science, according to which it proceeds by making bold conjectures, and then attempting to falsify those conjectures. In the simplest version of this account, when a hypothesis makes a prediction which is found to be false in an experiment, the hypothesis is rejected as falsified. The logic of this procedure is fully deductive. The hypothesis entails the prediction, and the falsity of the prediction refutes the hypothesis by *modus tollens*. Thus, Popper claimed that science was not based on the extrapolative inferences considered by Hume. The consequence then is that it is not so important, at least for science, if those inferences would lack a rational foundation.

Popper's account appears to be incomplete in an important way. There are always many hypotheses which have not yet been refuted by the evidence, and these may contradict one another. According to the strictly deductive framework, since none are yet falsified, they are all on an equal footing. Yet, scientists will typically want to say that one is better supported by the evidence than the others. We seem to need more than just deductive reasoning to support practical decision-making (Salmon 1981). Popper did indeed appeal to a notion of one hypothesis being better or worse "corroborated" by the evidence. But arguably, this took him away from a strictly deductive view of science. It appears doubtful then that pure deductivism can give an adequate account of scientific method.

Bibliography

- Achinstein, Peter, 2010, "The War on Induction: Whewell Takes on Newton and Mill (Norton Takes on Everyone)", *Philosophy of Science*, 77(5): 728–739.
- Armstrong, David M., 1983, *What is a Law of Nature?*, Cambridge: Cambridge University Press.
- Baier, Annette C., 2009, *A Progress of Sentiments*, Harvard: Harvard University Press.
- Bayes, Thomas, 1764, "An Essay Towards Solving a Problem in the Doctrine of Chances", *Philosophical Transactions of the Royal Society of London*, 53: 370–418.
- Beauchamp, Tom L, and Alexander Rosenberg, 1981, *Hume and the Problem of Causation*, Oxford: Oxford University Press.
- Bertrand, Joseph Louis Francois, 1888, *Calcul des probabilités*, Paris: Gauthier-Villars.
- BonJour, Laurence, 1998, *In Defense of Pure Reason: A Rationalist Account of A Priori Justification*, Cambridge: Cambridge University Press.
- Borel, Emile, 1909, *Elements de la theorie des probabilités*, Paris: Herman et Fils.

- Brown, M.B., 1987, "Review of *The Rationality of Induction*, D.C. Stove [1986]", *History and Philosophy of Logic*, 8(1): 116–120.
- Burks, Arthur W., 1953, "The Presupposition Theory of Induction", *Philosophy of Science*, 20(3): 177–197.
- , 1955, "On the Presuppositions of Induction", *Review of Metaphysics*, 8(4): 574–611.
- Campbell, Scott, 2001, "Fixing a Hole in the Ground of Induction", *Australasian Journal of Philosophy*, 79(4): 553–563.
- Campbell, Scott, and James Franklin, 2004, "Randomness and the Justification of Induction", *Synthese*, 138(1): 79–99.
- Carnap, Rudolph, 1950, *Logical Foundations of Probability*, Chicago: University of Chicago Press.
- , 1952, *The Continuum of Inductive Methods*, Chicago: University of Chicago Press.
- Carroll, John W., 2016, "Laws of Nature", *Stanford Encyclopedia of Philosophy* (Fall 2016 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/fall2016/entries/laws-of-nature/>>.
- Carroll, Lewis, 1895, "What the Tortoise said to Achilles", *Mind*, 4(14): 278–280.
- Cesa-Bianchi, Nicolo, and Gabor Lugosi, 2006, *Prediction, Learning, and Games*, Cambridge: Cambridge University Press.
- Chart, David, 2000, "Schulte and Goodman's Riddle", *British Journal for the Philosophy of Science*, 51(1): 147–149.
- Cleve, James van, 1984, "Reliability, Justification, and the Problem of Induction", *Midwest Studies In Philosophy*: 555–567.
- Cox, R. T., 1946, "Probability, frequency and reasonable expectation", *American Journal of Physics*, 14: 1–10.
- , 1961, *The Algebra of Probable Inference*, Baltimore, MD: Johns Hopkins University Press.
- de Finetti, Bruno, 1964, "Foresight: its logical laws, its subjective sources", in H.E. Kyburg (ed.), *Studies in subjective probability*, New York: Wiley, pp. 93–158.
- de Pierris, Graciela and Michael Friedman, 2013, "Kant and Hume on Causality", *The Stanford Encyclopedia of Philosophy* (Winter 2013 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/win2013/entries/kant-hume-causality/>>.
- Dretske, Fred I., 1977, "Laws of Nature", *Philosophy of Science*, 44(2): 248–68.
- Eckhardt, Arnold, 2010, "Can the Best-Alternative-Justification Solve Hume's Problem? (On the limits of a promising approach)", *Philosophy of Science*, 77(4): 584–593.
- Feigl, Herbert, 1950, "De Principiis non disputandum", in Max Black (ed.), *Philosophical Analysis*, Ithaca, NY: Cornell University Press, pp. 119–56.
- Foster, John, 2004, *The Divine Lawmaker: Lectures on Induction, Laws of Nature and the Existence of God*, Oxford: Clarendon Press.
- Garrett, Don, 2002, *Cognition and Commitment in Hume's Philosophy*, Oxford: Oxford University Press.
- Ghosal, S., J. K. Ghosh, and A.W. van der Vaart, 2000, "Convergence rates of posterior distributions", *The Annals of Statistics*, 28: 500–531.
- Ghosal, S., J. Lember, and A. W. van der Vaart, 2008, "Non-parametric Bayesian model selection and averaging", *Electronic Journal of Statistics*, 2: 63–89.

- Giaquinto, Marcus, 1987, “Review of *The Rationality of Induction*, D.C. Stove [1986]”, *Philosophy of Science*, 54(4): 612–615.
- Goodman, Nelson, 1955, *Fact, Fiction and Forecast*, Cambridge, MA: Harvard University Press.
- Hacking, Ian, 1975, *The Emergence of Probability: a Philosophical Study of Early Ideas About Probability, Induction and Statistical Inference*, Cambridge: Cambridge University Press.
- Hájek, Alan, 2011, “Interpretations of Probability”, *The Stanford Encyclopedia of Philosophy* (Winter 2012 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/win2012/entries/probability-interpret/>.
- Harman, Gilbert, 1968, “Enumerative Induction as Inference to the Best Explanation”, *Journal of Philosophy*, 65(18): 529–533.
- Henderson, Leah, 2014, “Bayesianism and Inference to the Best Explanation”, *The British Journal for the Philosophy of Science*, 65(4): 687–715.
- Howson, Colin, 2000, *Hume’s Problem: Induction and the Justification of Belief*, Oxford: Oxford University Press.
- , 2011, “No Answer to Hume”, *International Studies in the Philosophy of Science*, 25(3): 279–284.
- Huemer, Michael, 2009, “Explanationist Aid for the Theory of Inductive Logic”, *The British Journal for the Philosophy of Science*, 60(2): 345–375.
- [T] Hume, David, 1739, *A Treatise of Human Nature*, Oxford: Oxford University Press. (Cited by book.part.section.paragraph.)
- [E] —, 1748, *An Enquiry Concerning Human Understanding*, Oxford: Oxford University Press. (Cited by section.part.paragraph.)
- Jackson, Alexander, 2019, “How to solve Hume’s problem of induction”, *Episteme* 16: 157–174.
- Jeffreys, Harold, 1939, *Theory of Probability*, Oxford: Oxford University Press.
- Johnson, William Ernest, 1921, *Logic*, Cambridge: Cambridge University Press.
- , 1932, “Probability: the Deductive and Inductive Problems”, *Mind*, 49(164): 409–423.
- Kant, Immanuel, 1781, *Kritik der reinen Vernunft*. Translated as *Critique of Pure Reason*, Paul Guyer and Allen W. Wood, A., (eds.), Cambridge: Cambridge University Press, 1998.
- , 1783, *Prolegomena zu einer jeden künftigen Metaphysik, die als Wissenschaft wird auftreten können*. Translated as *Prolegomena to Any Future Metaphysics*, James W. Ellington (trans.), Indianapolis: Hackett publishing, 2002.
- Kelly, Kevin T., 1996, *The Logic of Reliable Inquiry*, Oxford: Oxford University Press.
- , 2007, “A new solution to the puzzle of simplicity”, *Philosophy of Science*, 74: 561–573.
- Kelly, Thomas, 2010, “Hume, Norton and induction without rules”, *Philosophy of Science*, 77: 754–764.
- Keynes, John Maynard, 1921, *A Treatise on Probability*, London: Macmillan.
- Lange, Marc, 2011, “Hume and the Problem of induction”, in Dov Gabbay, Stephan Hartmann and John Woods (eds.), *Inductive Logic, (Handbook of the History of Logic, Volume 10)*, Amsterdam: Elsevier, pp. 43–92.
- Laplace, Pierre-Simon, 1814, *Essai philosophique sur les probabilités*, Paris. Translated in 1902 from the sixth French edition as *A Philosophical Essay on Probabilities*, by Frederick Wilson Truscott and Frederick Lincoln Emory, New York: John Wiley and Sons. Retranslated in 1995 from the fifth

- French edition (1825) as *Philosophical Essay on Probabilities*, by Andrew I. Dale, 1995, New York: Springer-Verlag.
- Maher, Patrick, 1996, “The Hole in the Ground of Induction”, *Australasian Journal of Philosophy*, 74(3): 423–432.
- Maxwell, Nicholas, 2017, *Understanding Scientific Progress: Aim-Oriented Empiricism*, St. Paul: Paragon House.
- Mitchell, Tom, 1997, *Machine Learning*: McGraw-Hill.
- Morris, William E., and Charlotte R. Brown, 2014 [2017], “David Hume”, *The Stanford Encyclopedia of Philosophy* (Spring 2017 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2017/entries/hume/>.
- Norton, John D., 2003, “A Material Theory of Induction”, *Philosophy of Science*, 70(4): 647–670.
- , 2010, “There are no universal rules for induction”, *Philosophy of Science*, 77: 765–777.
- , 2021, *The Material Theory of Induction*: BPS Open/University of Calgary Press.
- Okasha, Samir, 2001, “What did Hume Really Show about Induction?”, *The Philosophical Quarterly*, 51(204): 307–327.
- , 2005a, “Bayesianism and the Traditional Problem of Induction”, *Croatian Journal of Philosophy*, 5(14): 181–194.
- , 2005b, “Does Hume’s Argument against Induction Rest on a Quantifier-Shift Fallacy?”, *Proceedings of the Aristotelian Society*, 105: 237–255.
- Owen, David, 1999, *Hume’s Reason*, Oxford: Oxford University Press.
- Papineau, David, 1992, “Reliabilism, Induction and Scepticism”, *The Philosophical Quarterly*, 42(166): 1–20.
- Popper, Karl, 1935 [1959], *Logik der Forschung*, Wien: J. Springer. Translated by Popper as *The Logic of Scientific Discovery*, London: Hutchinson, 1959.
- Ramsey, Frank P., 1926, “Truth and Probability”, in R.B. Braithwaite (ed.), *The Foundations of Mathematics and Other Logical Essays*, London: Routledge and Kegan-Paul Ltd., pp. 156–98.
- Reichenbach, Hans, 1949, *The Theory of Probability*, Berkeley: University of California Press.
- , 1938 [2006], *Experience and Prediction: An Analysis of the Foundations and the Structure of Knowledge*, Chicago: University of Chicago Press. Page numbers from the 2006 edition, Indiana: University of Notre Dame Press.
- Romeijn, Jan-Willem, 2004, “Hypotheses and Inductive Predictions”, *Synthese*, 141(3): 333–364.
- Russell, Bertrand, 1946, *A History of Western Philosophy*, London: George Allen and Unwin Ltd.
- , 1948, *Human Knowledge: Its Scope and Limits*, New York: Simon and Schuster.
- Salmon, Wesley C., 1963, “On Vindicating Induction”, *Philosophy of Science*, 30(3): 252–261.
- , 1966, *The Foundations of Scientific Inference*, Pittsburgh: University of Pittsburgh Press.
- , 1981, “Rational Prediction”, *British Journal for the Philosophy of Science*, 32(2): 115–125.
- Salmon, Wesley C., 1953, “The Uniformity of Nature”, *Philosophy and Phenomenological Research*, 14(1): 39–48.
- Savage, Leonard J., 1954, *The Foundations of Statistics*, New York: Dover Publications.

- Schulte, Oliver, 1999, “Means-Ends Epistemology”, *British Journal for the Philosophy of Science*, 50(1): 1–31.
- , 2000, “What to believe and what to take seriously: a reply to David Chart concerning the riddle of induction”, *British Journal for the Philosophy of Science*, 51: 151–153.
- , 2017 [2018], “Formal Learning Theory”, *The Stanford Encyclopedia of Philosophy* (Spring 2018 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2018/entries/learning-formal/>.
- Schurz, Gerhard, 2008, “The Meta-inductivist’s Winning Strategy in the Prediction Game: A New Approach to Hume’s Problem”, *Philosophy of Science*, 75(3): 278–305.
- , 2017, “Optimality Justifications: New Foundations for Foundation-Oriented Epistemology”, *Synthese*, 73:1–23.
- , 2019, *Hume’s Problem Solved: the Optimality of Meta-induction*, Cambridge, MA: MIT Press.
- , 2021a, “Meta-induction over unboundedly many prediction methods: a reply to Arnold and Sterkenburg”, *Philosophy of Science*, 88: 320–340.
- , 2021b, “The No Free Lunch Theorem: bad news for (White’s account of) the problem of induction”, *Episteme*, 18: 31–45.
- Schurz, Gerhard, and Paul Thorn, 2020, “The material theory of object-induction and the universal optimality of meta-induction: two complementary accounts”, *Studies in History and Philosophy of Science A*, 82: 99–93.
- Skyrms, Brian 2000, *Choice and Chance: an introduction to inductive logic*, Wadsworth.
- , 2012, *From Zeno to Arbitrage: Essays on Quantity, Coherence and Induction*, Oxford: Oxford University Press.
- Sober, Elliott, 1988, *Reconstructing the Past: Parsimony, Evolution and Inference*, Cambridge MA: MIT Press.
- Steel, Daniel, 2010, “What If the Principle of Induction Is Normative? Formal Learning Theory and Hume’s Problem”, *International Studies in the Philosophy of Science*, 24(2): 171–185.
- Sterkenburg, Tom, 2019, “The meta-inductive justification of induction: the pool of strategies”, *Philosophy of Science*, 86: 981–992.
- , 2020, “The meta-inductive justification of induction”, *Episteme*, 17: 519–541.
- , forthcoming, “Explaining the success of induction”, *British Journal for the Philosophy of Science*, <https://doi.org/10.1086/717068>.
- Sterkenburg, Tom and Peter Grünwald, 2021, “The no-free-lunch theorems of supervised learning”, *Synthese*, 199: 9979–10015.
- Stove, David C., 1986, *The Rationality of Induction*, Oxford: Clarendon Press.
- Strawson, Peter Frederick, 1952, *Introduction to Logical Theory*, London: Methuen.
- Tooley, Michael, 1977, “The Nature of Laws”, *Canadian Journal of Philosophy*, 7(4): 667–698.
- White, Roger, 2015, “The problem of the problem of induction”, *Episteme*, 12: 275–290.
- Will, Frederick L., 1948, “Donald Williams’ Theory of Induction”, *Philosophical Review*, 57(3): 231–247.
- Williams, Donald C., 1947, *The Ground of Induction*, Harvard: Harvard University Press.

- Wittgenstein, Ludwig, 1953, *Philosophical Investigations*, New Jersey: Prentice Hall.
- Wolpert, D. H., 1997, “No free lunch theorems for optimization”, *IEEE Transactions on Evolutionary Computation*, 1: 67–82.
- , 1992, “On the connecton between in-sample testing and generalization error”, *Complex Systems*, 6: 47–94.
- , 1996, “The lack of a priori distinctions between learning algorithms”, *Neural Computation* 8: 1341–1390.
- Worrall, John, 2010, “For Universal Rules, Against Induction”, *Philosophy of Science*, 77(5): 740–53.
- Wright, Crispin, 2004, “Wittgensteinian Certainties”, in Denis McManus (ed.), *Wittgenstein and Scepticism*, London: Routledge, pp. 22–55.
- Zabell, Sandy L., 1988, “Symmetry and Its Discontents”, in Brian Skyrms (ed.), *Causation, Chance and Credence*, Dordrecht: Springer Netherlands, pp. 155–190.
- , 1989, “The Rule of Succession”, *Erkenntnis*, 31(2–3): 283–321.

Academic Tools

 [How to cite this entry.](#)

 [Preview the PDF version of this entry at the Friends of the SEP Society.](#)

 [Look up topics and thinkers related to this entry at the Internet Philosophy Ontology Project \(InPhO\).](#)

 [Enhanced bibliography for this entry at PhilPapers](#), with links to its database.

Other Internet Resources

- Vickers, John, “The Problem of Induction,” *Stanford Encyclopedia of Philosophy* (Spring 2018 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2018/entries/induction-problem/>. [This was the previous entry on the problem of induction in the *Stanford Encyclopedia of Philosophy* — see the [version history](#).]
- [Teaching Theory of Knowledge: Probability and Induction](#), organization of topics and bibliography by Brad Armendt (Arizona State University) and Martin Curd (Purdue).
- [Forecasting Principles](#), A brief survey of prediction markets.

Related Entries

[Bayes' Theorem](#) | [belief, formal representations of](#) | [confirmation](#) | [epistemology, formal](#) | [Feigl, Herbert](#) | [Goodman, Nelson](#) | [Hume, David](#) | [Kant, Immanuel: and Hume on causality](#) | [laws of nature](#) | [learning theory, formal](#) | [logic: inductive](#) | [Popper, Karl](#) | [probability, interpretations of](#) | [Reichenbach, Hans](#) | [simplicity](#) | [skepticism](#) | [statistics, philosophy of](#) | [Strawson, Peter Frederick](#)

Acknowledgments

Particular thanks are due to Don Garrett and Tom Sterkenburg for helpful feedback on a draft of this entry. Thanks also to David Atkinson, Simon Friederich, Jeanne Peijnenburg, Theo Kuipers and Jan-

Willem Romeijn for comments.

Copyright © 2022 by
Leah Henderson <l.henderson@rug.nl>