Varieties of Justification—How (Not) to Solve the Problem of Induction

Marius Backmann¹

Received: 14 May 2018 / Accepted: 15 October 2018 / Published online: 8 November 2018
© The Author(s) 2018

1 Introduction

The debate about induction is a mess, perhaps only surpassed in its messiness by the debate about free will. There are almost as many different proposed solutions to the problem of induction as there are different formulations of the problem itself. While there seems to be a rather broad consensus that the problem is insoluble, the standards for what would count as a justification, were one available, vary wildly. That is not to say that the different authors who formulated the problem thought that a solution was possible, but even in a sceptic argument designed to show that a solution is unobtainable, we can see what a solution would need to entail. Obviously, what could count as a justification of induction depends on how one chooses to formulate the problem. In this paper, I will differentiate between three different standards for a possible justification. The three standards for a justification of induction are (1) to demonstrate how valid inductive inferences can be truth-preserving, (2) to demonstrate how induction can be truth-conducive, and (3) to show that inductive practice is rational. I will argue that the first two are unavailable, whereas the latter is at least in principle obtainable, although I will not argue for a particular proposal here.

The very first thing we need to make clear in this paper before we can even start the investigation into the varieties of justification is to make explicit what sort of inferences are affected by the classical problem of induction. Following this, I will give a brief survey of prominent formulations of the problem. We will then turn to the three abovementioned standards for a justification of induction. I will argue that the first two standards are impossible to meet, because they inevitably lead into some version of Hume's classic dilemma. This is the case because they all require the condition that nature will remain regular, which can neither be known *a priori*, nor, without entering a vicious circle, by induction. The only standard that is at least possible to satisfy

✉ Marius Backmann
marius.backmann@uni-konstanz.de

¹ Universität Konstanz, Konstanz, Germany

is to demonstrate that inductive practice is rational. Although rarely put forward as an attempt to solve the problem of induction, there exist a number of arguments for the rationality of conditionalisation in formal epistemology. If one accepts that inductive reasoning is a form of conditionalisation, for which I will briefly argue, then any argument for the rationality of conditionalisation is an argument for the rationality of inductive practice. The two distinct arguments I will portray here are (1) that only conditionalisers are immune to diachronic Dutch books, or Dutch strategies, and (2) that conditionalisation maximises expected epistemic utility. I will argue that these arguments do at least not necessarily fall prey to Hume's dilemma. Accordingly, this is the only way forward for a possible justification of induction.

2 Inductive Inferences

Induction is surprisingly hard to define, and as we will see, the classic characterisations of induction as an inference from particular premises to a general conclusion, or, even more specifically, from particular observations to a general law, are much too narrow. In the following, I will treat inductive inferences as a subset of ampliative inferences. An inference is ampliative if the content of the conclusion goes beyond the content of the premises. Whereas defining induction as an inference from particularity to generality yields a too narrow notion, simply equating inductive and ampliative inferences yields a definition that is too wide. There are more forms of ampliative inferences other than induction that are in need of justification. The most prominent form of inferences of this kind is the inference to the best explanation (IBE). Like induction, IBE is ampliative because the content of the conclusion is not contained in the content of the premises.¹

In the following, we will focus only on enumerative induction. Inferences to the best explanation (IBE), while they are clearly ampliative inferences, are structurally different from inductive inferences. In an IBE, we infer the conclusion via a bridge principle that claims that the premise is the best explanation for the conclusion. In order to demonstrate the difference, let us take a look at the master argument for scientific realism, which can be interpreted as an IBE:

- P₁: Science is successful.
- P₂: The success of science would be a miracle if our best scientific theories were not at least approximately true.
- P₃: That our current best scientific theories are at least approximately true is the best explanation for the success of science.
- C: Our current best scientific theories are at least approximately true.

Here, the conclusion is inferred via an implicit bridge principle that the explanatory virtue of the explanans is truth-conducive. In classical induction, the explanatory virtue of the conclusion plays no role. In its simplest form, induction is merely an

¹Lipton (2004), esp. 55-64

inference that a certain characteristic of a sample will be retained in different samples of the same population or in the population in general. The most obvious form of this inference is classical enumeration, which both Hume's and Popper's famous expositions of the problem of induction were concerned with. Enumerative induction is traditionally taken to be an inference from a number of particular instances to a generalisation. It is an inference of the pattern:

P₁: *a* is an *F* and a *G*.

P₂: *b* is an *F* and a *G*.

P_{*i*}: *i* is an *F* and a *G*, etc.

C: all *F*s are *G*s.

However, enumerative induction from particularity to generality is not the only form of inductive inference. It is for example possible to inductively infer a particular conclusion from a general premise, as long as both concern the distribution of the same trait in subsets of the same population, such as the following:

P₁: All observed *F*s have been *G*s.

C: The next *F* we will observe will also be a *G*.

Apart from inferences from inferences about the properties of particulars, inferences about the proportional distribution of certain traits in a population also seem to be structurally equivalent to the inductive inferences we discussed so far, and should also be treated as a subclass of inductive inferences:

P₁ Of *n* observed *F*s, *m/n* have been *G*s.

C *m/n* of all *F*s are *G*s.

What unites all of these different types of inferences is that they are ampliative inferences from the nature of a subset of a population to the nature of a different subset of the same population, or the population in general. In these inferences, in contrast to IBE, where we are supposed to infer the conclusion because of its explanatory virtue, we infer the conclusion merely on the basis of similarity between what we perceive a large enough sample of a population. Hence, I will treat all these inferences as instances of enumerative induction, even if they are inferences from general premises to a particular conclusion.

The problem of finding a justification for the ampliative step in inductive inferences is what constitutes the old problem of induction. These inferences are ampliative since the content of the conclusion goes beyond the content of the premises. In these inferences, we are referring to different, or more, instances in the conclusion than we are talking about in the premises (*all* ravens, as opposed to all observed ravens e.g.); hence, we cannot rely on the inference to be truth-preserving. Unlike in deductive inferences,² the conclusion does not need to be true if the premises are.

²Interestingly, the case for a justification of deductive inferences seems problematic as well: all arguments for why deduction is justifiable are themselves deductive, which renders any attempt to justify deduction circular (see, e.g. Haack 1976)

But if induction is not truth-preserving, why are justified to use it? Let us now turn to the archetypical formulations of the old riddle and their accompanying standards for a possible solution.

3 The Old Riddle(s)

David Hume gave the maybe still best-known exposition of the problem of induction in the *Treatise of Human Nature*.³ Famously, Hume formulated the old riddle as a dilemma, which serves as a template for a lot of recent formulations of the problem. Hume argues that if we wanted to demonstrate that inductive inferences were, in his words, “a product of reason”, that is if we wanted to demonstrate the validity of inductive inferences, we needed to know that nature is regular, or uniform. Only through this could we be justified to infer the conclusion of an enumerative argument. Hume gives no clear account of what exactly the extra assumption that nature is uniform is supposed to do for the respective inductive inferences. Most straightforwardly, we could add it as an extra premise to any enumerative inference. By adding this further premise, the enumerative inference is transformed into a valid deductive argument. We will discuss below why this is a problematic move. A classical enumerative inference would then look something like this in Hume’s analysis:

- P₁ All observed ferromagnets have so far attracted iron.
- P₂ Nature is uniform.
- P₃ If nature is uniform, then ferromagnets do not change their behaviour.
- C All future and unobserved ferromagnets attract iron.

However, *that* nature is uniform is a premise, which we again would need to justify. According to Hume, there are two options available for this: either we formulate an *a priori*, deductive (or, in Hume’s words, “demonstrative” inference with the conclusion that nature is uniform, or we formulate an *a posteriori*, inductive (or “probabilistic”) inference to that effect. The latter strategy would amount in a circularity, because we needed an inductive inference to justify the premise that would in turn justify inductive inferences in general. The first horn of the dilemma is slightly harder to analyse. That the uniformity of nature is arrived at by a deductive inference does not seem problematic at first glance. The clue is that Hume holds that this would make the uniformity of nature a necessary fact. He seems to hold that if we could give a deductive justification for the uniformity of nature, this would not be an empirical inference, one whose premises are arrived at by observation. Rather, it would make the uniformity of nature an *a priori* truth. But that nature is uniform seems like an empirical fact, one that we cannot justify *a priori*. So the problem that the first horn of Hume’s dilemma expresses is not that the uniformity of nature could be an *a priori* fact, but that *a priori* knowledge about whether nature is uniform or not is

³Hume (2000)

unavailable. We simply cannot know whether nature is uniform without empirically investigating what nature is like.

The other *locus classicus* of the debate about the old riddle is Popper's famous exposition in the beginning of *The Logic of Scientific Discovery*.⁴ Since Popper's way of stating the problem is relevant for my argument, let's take a look at the exact formulation:

The problem of induction may also be formulated as the question of the validity or the truth of universal statements which are based on experience [...]. [...] Accordingly, people who say of a universal statement that we know its truth from experience usually mean that the truth of this universal statement can somehow be reduced to the truth of singular ones, and that these singular ones are known by experience to be true; which amounts to saying that the universal statement is based on inductive inference.⁵

In the following, Popper, like Hume, formulates the problem the validity of inductive inferences as a dilemma. What we needed to justify induction is to find a "principle of induction", that is a principle that would demonstrate the logical validity of inductive inferences. The goal for a justification seems to be here to make inductive inferences more like deductive inferences such a way that "[...] the truth of [a] universal statement can somehow be reduced to the truth of singular ones [...]". This could be achieved by a set of logical norms that would demonstrate the logical validity of inductive inferences. Popper claims that the principle of induction that is supposed to demonstrate the validity of inductive inferences cannot be a "tautology" like the rules of classical two-valued deductive logic. If the principle of induction were tautological, there would not be a problem of induction. The reason Popper gives why we cannot arrive at this principle of induction is a version of Hume's dilemma: either we knew *a priori* that induction were justifiable or we would have to obtain this principle via experience. Again, *a priori* knowledge of the validity of inductive inferences is impossible for the same reasons as above, and an empirical justification would involve an inductive inference from past success of induction to its future success and would hence be circular.⁶ Both Popper's and Hume's dilemma sought to demonstrate that inductive inferences, in order to be justifiable, would be furnished with an additional premise, which is unavailable because of the respective dilemma. Remember that according to Hume, a justifiable inductive inference would take approximately this form:

- P₁: All observed ferromagnets have so far attracted iron.
- P₂: Nature is uniform.
- P₃: If nature is uniform, then ferromagnets do not change their behaviour.
- C: All future and unobserved ferromagnets attract iron.

⁴Popper (2005), 4-5

⁵Karl (1966), 4.

⁶See Popper (2005), 4-5

For Popper, a justifiable inductive inference would rather have the following form:

- P₁: All observed ferromagnets have so far attracted iron.
 P₂: (Principle of induction:) From past or observed instances infer to future or unobserved instances.
 C: All future and unobserved ferromagnets attract iron.

In both cases, premise 2) is unavailable because of the respective dilemma. And crucially, both inferences are being turned into deductive ones by the addition of an otherwise suppressed premise (premise 2) in both cases). It seems that according to Popper's and Hume's way of framing the problem of induction, the problem could be solved if we *eliminated* inductive reasoning and reduced it to truth-preserving reasoning. But before we turn to the general issue with that way of framing the question, let us briefly take a look at the content of these additional premises.

3.1 The Uniformity of Nature

Popper never gives an account for what a principle of induction would actually look like, but supposedly it looks something like I sketched above: (PI) *From past or observed instances infer to future or unobserved instances*. This principle, like the suppressed premise outlined in Hume's exposition of the problem, relies on the uniformity of nature. If nature were irregular and could behave entirely different tomorrow than how it behaved so far, (PI) would be a poor guide to reasoning. We would only be justified to infer the unobserved from the observed (or the future from the past and present) if we knew that the observed were any guide to what the unobserved looked like. That would not be the case if nature were irregular. So both Popper's and Hume's account for what would be needed to justify induction, if that were possible, rely on the uniformity of nature.

At this point, we should briefly clarify what the uniformity, or regularity, of nature is supposed to be. Clearly, uniformity is not to be confused with determinism. The success of a justification of induction according to this standard is not dependent on us being able to affirm in a non-circular manner that the world is *deterministic*. A world could be indeterministic if, e.g., there were stochastic laws, but it would still be uniform in the sense that these stochastic regularities remained stable. To take an overused but helpful example, radioactive decay is an indeterministic process. Neither can we predict for every specific radioactive particle when it is going to decay, nor is there, if radioactive decay is indeed ontologically indeterministic, a hidden and undetectable factor that determines when any given radioactive particle is going to decay. However, the half-life of a radioactive isotope is stable. The half-life of uranium-238, for example, is 4.468 billion years. That means that within 4.468 billion years, approximately half of any sample of uranium-238 will have decayed. But although radioactive decay is an indeterministic process, a world (like ours) which is indeterministic in that way can still be regular, or uniform. The half-life of uranium-238 does not change: it is not 4.468 billion years for the first 10 billion years, and

then switches to 4 years for the next 10 billion years, and so on. That is to say that the indeterministic regularities stay in place in a uniform world.⁷

Unfortunately, however, that a world is uniform does not entail that the *observed* regularities remain stable. They could change wildly, which could be the consequence of undetected, and undetectable, fundamental regularities, which could remain stable. Such a world would appear irregular, but it would not be. So even if we knew that nature is uniform, and even if knew that *a priori*, that still would not entail that the observed regularities would remain stable. Their changing, if it is the consequence of more fundamental, but (in principle) undetectable regularities, would still not imply that nature is irregular. But then, the regularity of nature would not help us if we wanted to make any inferences about our observed regularities truth-preserving. Hence, even in a regular world, our inductive inferences would be defeasible. We would even need a second suppressed premise apart from that nature is regular, namely that we are able to tell which apparent irregularities are consequences of more fundamental, but stable regularities.

3.2 The Issue with the Classical Formulations of the Old Riddle

As we have seen in the reconstruction of what a justifiable inference according to Hume and Popper would look like, the justifiable inference would in both accounts cease to be inductive. This suggests that Hume and Popper both hold that inductive inferences could only be justified if they could be treated as enthymemes of deductive inferences. Enthymemes are inferences with a suppressed premise.⁸ By turning inductive inferences into enthymemes of deductive inferences, they would gain all the properties inductive inferences essentially lack, which is why they were in need of justification in the first place: they would cease to be ampliative and become truth-preserving, that is we could, like in deductive inferences, be certain that the conclusion is true if the premises are true. That is what Popper presumably meant when he said that in order to solve the problem, we would need to show that “the truth of [a] universal statement can somehow be reduced to the truth of singular ones, and that these singular ones are known by experience to be true; [...]”⁹

However, since inductive inferences, if we treat them as enthymematic inferences, would be turned into deductive inferences, we would not *justify* induction, but *eliminate* inductive reasoning. In a way, the problem of induction in this view is not that we fail to justify inductive inferences, but that we fail to turn them into deductions, because we cannot justify the suppressed premise of the enthymematic inferences. These classical formulations of the problem entail that a justification is generally

⁷ At this stage, it should be noted that the New Riddle rears its head. Whether the world is regular is not only dependent on the world, but also on our description of the world. While a world in which emeralds change their colour might appear irregular in our language, it appears regular if we use a *grue/bleen* language. For the remainder of the paper, I will assume that the reference frame remains fixed. I will explore this issue in a different paper.

⁸ Laurence Bonjour also reconstructs both Hume’s and Popper’s accounts of the problem as the problem to justify the suppressed premise in an enthymematic argument (see, Bonjour (1998), 190).

⁹ Popper (2005), 4

unobtainable, just as Hume diagnosed: any attempt to justify induction this way implies that we can fill in the extra suppressed premises of an enthymematic inference. This requires exactly the knowledge we lack, which is the very reason we have to infer inductively at all: it requires complete knowledge of the phenomena we are reasoning about. To establish that nature is uniform, we would either enter a regress or would have to claim that the uniformity of nature is *a priori* knowable. Any attempt to make induction truth-preserving means to argue that justified induction is an enthymeme of a truth-preserving inference. There is no way to do this without falling into some version of Hume's dilemma.

This diagnosis is entirely unsurprising, given that Hume's and Popper's view was that the problem is insoluble. One could ask, however, whether their demands for what a justification should entail were not maybe unfair. To treat the problem of induction in this way would also entail that induction, as long as it is a justifiable sort of inference, is not *sui generis*: justifiable induction would be deduction in disguise. Granted, Hume and Popper held that the problem was insoluble, so there would be *sui generis* induction, but only because the inferences are unjustifiable and we fail to reduce these inferences to deductive ones. The problem of induction would turn out to be that there even *are* inductive inferences. Peter Strawson has come to a similar conclusion: the classical formulation of the problem of induction holds inductive inferences to a standard they cannot fulfil: that of deductive inferences. If we wanted to honour the ampliative nature of inductive inferences, we should instead look for a way to justify induction in a way that pays reference to its nature as genuinely ampliative and non-truth-preserving inferences.¹⁰

One way to find a justification that would pay respect to the *sui generis* nature of inductive reasoning would be to demand of a justification that it does not have to demonstrate how induction could be truth-preserving but truth-conducive. A possible justification would then take the form of merely showing that in an inductive inference, the truth of the premises at least raises the probability that the conclusion is true. In the following, I will discuss proposals to that effect and ultimately reject them.

4 Alternative Justification: Truth-Conduciveness and Reichenbach's Pragmatic Vindication

To require a justifiable inductive inference to be truth-preserving seems to put an unreasonably strong demand on inductive inferences and completely neglects their ampliative nature. After all, induction is by its very nature defeasible, so it is no surprise that any attempt to reduce it to an enthymematic deduction is bound to fail. However, it is less of an outrageous demand to want to know how, even in a defeasible inference, the premises can confer some degree of truth or likelihood of truth on the conclusion. It is this less demanding standard that has often been employed by people who are less sceptical about the prospect of justifying induction. As we shall

¹⁰Strawson (1952), 250.

see, it is the standard that stands behind the classical probabilistic attempts to solve the puzzle, the Stove-Williams account of induction, and Laurence Bonjour's view.¹¹ The general idea is that a justification should not demonstrate that inductive inferences preserve truth like an enthymematic deductive inference, but that the premises confer *some* degree, or likelihood, of truth upon the conclusion. As a point of departure, let us take Bonjour's account of what an epistemic justification of induction would entail:

If we understand epistemic justification [...] as justification that increases to some degree the likelihood that the justified belief is true and that is thus conducive to finding the truth, the issue is whether inductive reasoning confers any degree of epistemic justification, however small, on its conclusion.¹²

This is consistent with what we find in the various probabilistic attempts to justify induction. In this paper, I will not differentiate *en détail* between the various probabilistic accounts. What these accounts all have in common is that they try to demonstrate how the accumulation of evidence make it increasingly probable that the conclusion is true. Take, e.g. the Stove-Williams account. David Stove and Donald Williams both independently hold the view that the law of large numbers could help to justify the inference that a large sample will resemble the population. According to them, we are justified to infer that based on the law of large numbers, the proportion of *F*s that are *G*s in large sample of a finite population will resemble the proportion of *F*s that are *G*s in the population. If we randomly draw a large sample of a population, it is very likely that the sample will be representative. If we drew any logically possible large sample from a population, the majority of these samples will show a distribution of the trait we are interested in that falls within a very small margin around the original distribution in the sample.¹³ It is thus probable that the conclusion of an inductive inference from the composition of a large sample to the composition of the population is true.¹⁴

While these accounts do not demand that induction can be made truth-preserving by being reduced to enthymematic deductive inferences, they unfortunately all, and necessarily so, share one problem with the abovementioned sceptical accounts. Necessarily, any attempt that even merely claims that the truth of the conclusion of an inductive inference becomes more likely as more and more evidence is accumulated, relies upon the uniformity of nature. That a sample, even a large one, and even in the long run, resembles the population (of which some members exist in the future or are so far unobserved) cannot be assumed with certainty if it is possible that nature can completely change tomorrow—if it is possible that tomorrow, negatively charged particles attract each other, cats bark, and dogs meow. To put it in the terms above, that a large sample will resemble the population is only apparent for a population that is not going to change. Let's say we want to infer whether all future ravens will

¹¹ See, e.g. Reichenbach (1935), Stove (1986), Williams (1947), and Bonjour (1998)

¹² Bonjour (1998), 189.

¹³ See Smart (2013), 327.

¹⁴ See Williams (1947) and Stove (1986).

continue to be black, as were the past and present (observed) ones. The population in question is the entirety of ravens—past, present, and future. That the future ravens will resemble our sample is only given if we know that in the future, the regularities will not change. If that were the case, our sample of ravens that we have observed before would not have been random. It would have been a temporally ordered and restricted sample from a section of time when ravens were not white, so the sample would have been skewed. And there is no way to tell that our sample is not skewed, unless we knew that nature is uniform. But we cannot, given Hume's dilemma.

So, these attempts might be more realistic in the sense that they do not hold that in order to be justified, induction would have to be eliminated, which would deny that induction is a *sui generis* type of inference. Yet, such a more realistic justification can still in principle never be achieved. That nature is uniform, and that the observed regularities will remain stable as a consequence of this uniformity of nature, would still have to be independently inferred: either deductively or inductively. And that exactly is Hume's dilemma. So even granted that we could solve all the problems associated with these accounts, such as the problem of how to demonstrate that a specific inductive inference rule is in any way better than, e.g. counter-induction, Hume's dilemma still bites.

So without going much deeper into the various extant probabilistic accounts, we can reject them altogether. So is all lost then? If even the more modest demand that the premises merely raise the probability of the truth of the conclusion is impossible to be met, what hope can we have? If it is at all possible to solve the problem of induction, what we would need for this is an account that did not presuppose the regularity of nature. In the next section, we will see that the various accounts that aim to demonstrate the rationality of specific formal frameworks for inductive inferences often do not require that nature is uniform. Maybe we can derive a justification of inductive *practice* from this, a justification of why we indeed use induction, even if that does not amount to a demonstration that the conclusion of an inductive inference is true or likely to be true. Before we move on to the next section, where we take a closer look at such attempts to justify induction. Let us briefly discuss Reichenbach's pragmatic vindication of induction, which can be understood as a sort of intermediary step towards such an account. As we will see, Reichenbach's account comes with one important restriction of the views discussed in the next section that do not have and which poses a severe problem for his account.

At first sight, Reichenbach's pragmatic vindication of induction and Wesley Salmon's reformulation of it seem to fit the same mould as the law of large numbers views discussed above.¹⁵ Salmon claims that induction can never be shown to be truth-preserving, but what we can assert is that inductive inferences are supported by the evidence stated in the premises. That is to say that, any scepticism about the possibility to solve Hume's dilemma at all aside, there is a relation between the truth of the premises and the truth of the conclusion. Crucially, if any method is successful in extending our knowledge in this view, induction is. Reichenbach's and Salmon's principles of induction are quite close to the Stove-Williams view in that they explicitly

¹⁵See Reichenbach (1935) and Salmon (1974).

refer to the importance of the long run. Salmon, for instance, proposes the following inductive rule:

[G]iven m/n of observed A are B , [...] infer that the ‘long run’ frequency of B among A is m/n .¹⁶

The crucial difference between Reichenbach’s original account and the law of large numbers views discussed above is that Reichenbach holds that while the uniformity of nature cannot be simply assumed to justify inductive reasoning, he holds that induction is the only sort of ampliative reasoning that could be successful *if* nature were uniform, and that if nature is not kind enough to be uniform, then *no* mode of reasoning can be successful. It would hence be irrational if we would not engage in inductive reasoning.¹⁷ The pragmatic vindication thus lies not within a demonstration that induction actually is in practice truth-conducive, but that it is the only mode of reasoning that even has a hope of being truth-conducive if nature is kind enough to us. It would hence be irrational for not to use inductive reasoning.

I will not discuss Reichenbach’s account in greater detail here. The view is notable for holding a middle ground between attempts at justification which allude to the truth-conduciveness of inductive reasoning, which he seems to presume in case that nature is uniform, and attempts that allude to the rationality of inductive practice. There even exists an interesting analysis of Reichenbach’s vindication on the grounds of modern decision theory by Michael J. Shaffer.¹⁸ The reason I will not discuss Reichenbach’s account any further is that it still presupposes that induction will be truth-conducive *given that* nature is uniform in order to argue why we should reason inductively. I have argued above that this assumption is false: we cannot expect induction to be truth-conducive, even if nature turns out to be uniform. But without that condition, Reichenbach’s vindication is on shaky grounds. However, Reichenbach’s account is notable for shifting the focus towards the *rationality* of induction, away from the question whether induction is actually successful. So let us now turn to justifications of the rationality of inductive practice that do not make any assumption about the truth-conduciveness of induction, uniformity of nature notwithstanding.

5 Rationality

From now on, we will depart from the traditional debate about induction and focus on the justification of formal belief updating norms. There exist a number of attempts to justify conditionalisation that are designed to show that conditionalisation is rational under certain rationality constraints. The arguments to justify conditionalisation are usually derived from more established arguments for probabilism, i.e. the view that we should treat our beliefs in a way that fulfils the axioms of probability theory. These arguments are of interest to our discussion here since if one accepts that inductive

¹⁶Salmon (1974), 50.

¹⁷See, e.g., Reichenbach (1935), chapter 11.

¹⁸See Shaffer (2014).

reasoning is a species of conditionalisation, then any argument for the rationality of conditionalisation is an argument for the rationality of inductive practice. In the following, I will treat conditionalisation as any formal probabilistic framework to update one's beliefs in the light of evidence. Inductive inferences can thus be construed as belief updating procedures in which the available evidence represents a sample of a population we are updating our beliefs about. Even for the strictest cases of enumerative induction, conditionalisation can be seen as a way to formalise each iterative step of individual observation and the resulting adjustment of our beliefs. Since most of the arguments I discuss in this section are formulated in a Bayesian framework, I will not discuss alternative frameworks for belief updating such as Ranking Theory, e.g. although I have no reason to doubt that the arguments can be adapted to Ranking Theory.¹⁹

Take the following example to illustrate why I take conditionalisation to be suitable to formally represent enumerative induction as sketched above. Let's say it is the late nineties, and I have never consciously listened to a song by Radiohead (obviously, *Creep* doesn't count). Before I listen to one of their songs for the first time, I have some prior belief about whether I like their music, which we will express in terms of subjective probability $Pr_1(A)$. At this time, my subjective probability that I will like any music that is new to me is probably quite low, because I am at this time a self-important late teen who thinks they've figured out music, and as such I am unimpressed by default. I am then presented with evidence E in the shape of one of the songs on Radiohead's *OK Computer*. After I have received such evidence, my new, updated belief that I like Radiohead $Pr_2(A)$ should now be equal the initial conditional belief that I will like their music, given that I am presented with positive evidence of their brilliance $Pr_1(A|E)$. Since E consisted in a song that I liked and my initial credence that I would like new music was quite low, my new, updated credence that I like Radiohead $Pr_2(A)$ is now higher than $Pr_1(A)$.

We can immediately see how such a way of belief updating is compatible with how we sketched induction above. The inference is defeasible in the sense that it is always possible to gather new evidence that could lead me to reevaluate my credence regarding whether I like Radiohead or not. The inference is also ampliative in the sense that I extrapolate from a sample of Radiohead's *œuvre* in order to infer whether I like their catalogue. We should be clear here that not necessarily all instances of conditionalisation can be reconstructed as inductive reasoning, but that inductive reasoning can be reconstructed as an a subset of all instances of conditionalisation. If I want to check whether there really is a bottle of milk in my fridge, open the fridge, and conditionalise upon my prior credence about the contents of my fridge, then clearly this is a case of conditionalisation, but not of (enumerative) induction.

Given that we can understand inductive inferences as a subset of all instances of conditionalisation, any justification of that particular updating rule is a justification of inductive practice. Let us now take a look at three different strategies of demonstrating the rationality of conditionalisation, and see whether they can avoid Hume's

¹⁹See Spohn (2012) for a comprehensive account of ranking theory and Spohn (2016) for a ranking-theoretic account of enumerative induction.

dilemma, and whether this might lead the way to a justification of inductive practice that is not concerned with whether induction is truth-conducive, but whether it is rational.

5.1 Dutch Strategies

Dutch book arguments have traditionally been applied to argue for the rationality of probabilism. The basic idea is that if your belief system adheres to the axioms of probability theory, you are immune to accepting bets that guarantee a sure loss. In probability theory, the probabilities of all the possible outcomes of a certain situation should add up to 1. And if you assign subjective probabilities to possible outcomes, these too should add up to 1. The Dutch book argument now matches these subjective probabilities with odds of bets a bookie might sell you. If you have a certain degree of belief that p , say 0.8, you will accept to buy a bet for 80 cents that pays 1 dollar in case p comes about, and you should pay no more than 20 cents on a bet that pays 1 dollar in case that p does not come about. A person with incoherent belief states might now be susceptible to accept a set of bets that will result in a sure loss for the agent. If you have a degree of belief of .6 that p , and again .6 that not- p , a bookie can sell you a bet for 60 cent each that pays 1 dollar if p comes about, or 1 dollar if not- p comes about. The two bets together cost 1.20 dollars, but you can only receive 1 dollar, whatever the outcome. This set of bets is called a Dutch book. Basically, we can demonstrate that an agent has incoherent belief states if it is possible to construct a Dutch book against them. It has been argued that adhering to the axioms of probability theory makes you immune to being Dutch-booked, which is why it is rational to adhere to them, given you want to maximise utility.

Since we are dealing with arguments for the rationality of inductive reasoning, we should take a look at rationality arguments not for probabilism, but for probabilistic belief updating, for conditionalisation. In order to construct a Dutch book argument for conditionalisation, the argument needs to be made diachronic, since conditionalisation concerns not only how coherent your credences are at a certain point in time, but how you update your credences after you have gathered relevant evidence. An iterated set of bets over time that deliver a sure loss for the agent is called a Dutch strategy. A person adhering to conditionalisation should adjust their credences in the following way. Say at a certain time t_1 you have a degree of belief $Pr_1(A)$, and a degree of belief how likely A is given a certain event E occurs that we would take as evidence for A : $Pr_1(A|E)$. Now after E has either come about or not at t_2 , your new credence that A should be equal the conditional subjective probability that A given E you had at t_1 : $Pr_2(A)=Pr_1(A|E)$.

It has been demonstrated that if an agent violated this simple updating rule, but rather update their credences in a way where the updated subjective probability is, e.g. $Pr_2(A) < Pr_1(A|E)$, they are susceptible to accepting a Dutch strategy. A bookie could sell them the following set of bets, which are all fair in the eyes of the agent:²⁰

²⁰See Teller (1973) and Lewis (1999), cited here after (Vineberg et al. 2016)

Bet 1) for $Pr_1(A \wedge E)$

Receive \$1 for $(A \wedge E)$

Receive \$0 if not

Bet 2) for $Pr_1(A|E)Pr_1(\neg E)$

Receive \$ $Pr_1(A|E)$ if $\neg E$

Receive \$0 if not

Bet 3) for $[Pr_1(A|E) - Pr_2(A)] \cdot Pr_1(E)$

Receive \$ $Pr_1(A|E) - Pr_2(A)$ if E

Receive \$0 if not

If $\neg E$ is the case, the agent loses bet (1), wins (2), and loses (3), which, if you factor in the price of all bets, leads to a net loss of the price of the third bet: $[Pr_1(A|E) - Pr_2(A)] \cdot Pr_1(E)$, since bets (1) and (2) cancel each other out in case that $\neg E$. If E is the case, the bookie will *buy* a fourth bet on A from the agent:

Bet 4) for $Pr_2(A)$

Receive \$1 if A

Receive \$0 if not

Since the agent we are dealing with here violates conditionalisation in that they set the probability of $Pr_2(A)$ lower than, instead of equal $Pr_1(A|E)$, the agent will in case that E occurs suffer a net loss of again $[Pr_1(A|E) - Pr_2(A)] \cdot Pr_1(E)$. Hence, no matter if E or $\neg E$ is the case, the agent will lose $[Pr_1(A|E) - Pr_2(A)] \cdot Pr_1(E)$. An agent who did not violate conditionalisation would not be susceptible to a Dutch strategy. Hence, it is rational to adhere to conditionalisation.

Whether diachronic Dutch book arguments are a good strategy to justify probabilism is controversial, especially since they offer merely a pragmatic argument for an epistemic norm: manage your beliefs thus and you will not be exploitable.²¹ I will not attempt to settle the debate here. However, we can see that a possible attempt at a justification of induction by demonstrating the rationality of inductive practice by employing diachronic Dutch book arguments does not necessarily lead into Hume's dilemma in the way that the standards of justification discussed above do. The beauty of a Dutch strategy argument is that it does not depend on the outcome of the bets: whatever the outcome, the agent will make a sure loss, if they fail to adhere to probabilism. Granted, diachronic Dutch books *do* depend on the stability of the bet: in its classic form above the argument depends on the agent fixing a certain way to violate conditionalisation beforehand. If the agent decides to violate the rule in a different manner over time, the betting strategy above does not necessarily lead to a sure loss for the agent.²² The argument also depends in a very trivial sense on the continued existence and commitment to the set of bets of both bookie and agent. But, crucially, the argument does not depend on nature staying uniform in the sense that a particular outcome comes about, or that the regularities remain stable. Whatever

²¹ See Joyce (1998), section 2.

²² van Fraassen (1989), p. 174.

happens, as long as bookie and agent stay committed to their bets and updating strategy, the agent is susceptible to be Dutch-booked if they violate conditionalisation, and remains immune if they don't.

Hence, if we want to accept avoiding Dutch books as a pragmatic measure of the rationality of conditionalisation, this would be a possible justification of a certain formalised form of inductive practice. Since Dutch books are a controversial rationality measure, especially for epistemic rationality, let us briefly take a look at another argument for the *epistemic* rationality of conditionalisation, i.e. that conditionalisation maximises expected epistemic utility.²³

5.2 Expected Epistemic Utility

One way to provide a justification for epistemic norms is such as the conditionalisation to employ the methodology of expected utility theory.²⁴ Expected utility theory is the orthodox framework of decision theory. The expected epistemic utility argument are of great interest here because they offer an epistemic justification rather than a pragmatic one.

Very briefly put, according to expected utility theory, an agent's choice is supposedly rational if it is the one of which the agent expects the greatest value according to their preferences. If I wish to eat ice cream, I expect to fulfil this preference by going out and buying some ice cream, and hence it is rational for me to go out and buy ice cream. In expected epistemic utility theory, we substitute the usual items of decision theory with epistemic analoga in order to use the formal framework of expected utility theory: instead of judging the rationality of a decision on which action to take, we are talking about the rationality of 'choosing'²⁵ a credence function based on the available relevant evidence and our epistemic norms. We can call this an epistemic act. The utility we are trying to maximise is again purely epistemic: given that being in a belief state that is closest to the truth is desirable, an agent maximises expected epistemic utility if they adopt a credence function that we can expect to be as close to the truth as possible, or, to phrase it differently, of being as accurate as possible (or least inaccurate). So if we value being in a belief state that is as close to the truth as possible, we are rational to adopt an epistemic norm such as conditionalising if it maximises epistemic utility in such a way that we can expect it to furnish us with credences that are as close to the truth as possible.

To take an easy example of conditionalisation, consider that I am faced with choosing one and exactly one ball of ice cream from an ice cream parlour with a very

²³These are not the only extant justifications of conditionalisation. Recently, e.g. R.A. Briggs and Richard Pettigrew explored a way to justify conditionalisation by demonstrating that it is accuracy dominant (Briggs and Pettigrew [forthcoming](#)). For the sake of this paper, I will stick to Dutch strategies and expected epistemic utility as successful examples for pragmatic and non-pragmatic justifications of conditionalisation.

²⁴See, e.g. Greaves and Wallace (2006), Leitgeb and Pettigrew (2010a, b).

²⁵This does not need to be a deliberate, conscious choice in the way that we use the term in action theory of the free will debate (see Greaves and Wallace 2006).

limited selection of only four flavours, two fruity (F), and two non-fruity (N): vanilla (V), chocolate (C), strawberry (S), and lemon (L). My partner is tasked with predicting my choice of ice cream. Since she knows me very well, she sets her priors for me choosing any particular ice cream flavour as follows:

$$\begin{aligned}\Pr_1(V) &: 0.75; \\ \Pr_1(C) &: 0; \\ \Pr_1(S) &: 0.1; \\ \Pr_1(L) &: 0.15.\end{aligned}$$

Her priors for F and N should hence be $\Pr_1(F): 0.25$ and $\Pr_1(N): 0.75$. If for some reason she knew that I will chose a non-fruity flavour, her credences would be as follows:

$$\begin{aligned}\Pr_1(V|N) &: 1; \\ \Pr_1(C|N) &: 0; \\ \Pr_1(S|N) &: 0; \\ \Pr_1(L|N) &: 0.\end{aligned}$$

And if she knew that I will chose a fruity flavour, her priors for the individual flavours, if she were to keep the ratio my of my tendency to choose lemon over strawberry intact, would be:

$$\begin{aligned}\Pr_1(V|F) &: 0; \\ \Pr_1(C|F) &: 0; \\ \Pr_1(S|F) &: 0.4; \\ \Pr_1(L|F) &: 0.6.\end{aligned}$$

Suppose I make my choice, and the ice cream sales person, without revealing my precise choice, tells her that I have uncharacteristically chosen a fruity ice cream flavour. Knowing my tendency to choose lemon over strawberry, and being a rational epistemic agent, she adjusts her posterior credences by conditionalising such that she sets her updated credences by simply taking her priors for me choosing the individual flavours given that I have chosen a fruity flavour. This leads her to the following credences, given that F

$$\begin{aligned}\Pr_2(V) &= \Pr_1(V|F) = 0; \\ \Pr_2(C) &= \Pr_1(C|F) = 0; \\ \Pr_2(S) &= \Pr_1(S|F) = 0.4; \\ \Pr_2(L) &= \Pr_1(L|F) = 0.6.\end{aligned}$$

Any alternative updating rule would yield a different distribution of updated credences.

It can now be shown that an agent maximises epistemic utility if they adopt the credence function provided by conditionalisation, but not if they adopt any alternative updating rule, which would yield a different credence function. Adopting the formalism of expected utility theory, we can now go ahead and assign a utility $U(s, p)$, represented by a real number, to every possible pair of credence function p and possible state s of the world, where these states of the world are understood as the

outcomes we assigned our credences to. In our example here, any choice of ice cream flavour represents a different possible state of the world. If the agent values truth, they are going to assign a high utility to any pair, where they have a high credence in the state that actually obtains. If we add up all the pairs of credence functions and states of the world with their respective assigned utility, we yield the agent's utility function. A rational agent will, given their prior credences, and relative to a credence function P , 'choose' the credence function that gives them the highest expected utility. The 'choosing' of the respective credence function is depicted here as an epistemic act a , whose expected utility function we can represent as follows:

$EU^P(a) : \sum_{s \in S} p(s) \cdot U(s, (a, s))$. For the sake of not cluttering up the paper with proofs, we will not go into the details here. Greaves and Wallace now go on to show that every other updating rule will yield a lower expected utility than conditionalising does, given that the agent places a high utility on being correct (Greaves and Wallace 2006, 615).

So, although this justification of conditionalisation is epistemic in the sense that it purports that an updating rule will maximise expected epistemic utility if it gets us close to the truth, it does not rely on what the world actually looks like. The important question here is whether the expected epistemic utility argument falls prey to Hume's dilemma. I argue it does not, at least not necessarily. Here, the regularity of nature plays a different role than for options 1) and 2). Whereas there, the truth-conduciveness, or even truth-preserving nature of inductive inferences could only be asserted (if even at all) if nature is uniform, this is not the case for expected epistemic utility arguments. Here, it does not matter if nature is regular or not for it to be true that conditionalisation maximises *expected* epistemic utility. That it does is just a consequence of conditionalisation and expected utility theory. So, if we accept the expected epistemic utility argument as a justification of conditionalisation, then it holds regardless of whether the world is regular or not.

However, this does *not* imply that conditionalisation always gives you the best results. It is easy to construct a world in which conditionalisation produces worse predictions than other updating rules, because the regularities might change in a way the conditionaliser cannot foresee: consider a world which will last for exactly 5,000 years. For the first 4,999 years, all swans are white. Only in the last year, all swans will be black. Importantly, there is no underlying mechanism that could account for the change. The distribution of swan colour is also not indeterministic in the sense that there is an indeterministic law that governs that a certain percentage of swans are either colour, and we were just unlucky that all the white ones occurred in the first 4,999 years. The world in this example is hence not just indeterministic in the sense that it contains statistical regularities, but it is a sort of flip-flop world: the laws change suddenly and for no underlying reason. In the year 4,999, a conditionaliser will have a pretty high credence that all swans are white. In contrast, someone who does not conditionalise might, regardless of the evidence, always believe that not all swans are white. That person would end up making the correct prediction regarding the colour of swans in the year 5,000, whereas the conditionaliser fails to predict correctly.

We could even construct a world in which conditionalisation *always* delivers the wrong result if we want to make a prediction. There could be a particularly nasty

world in which every time somebody conditionalises over the evidence and their priors to make a prediction, the exact opposite of what was indicated in the agent's posterior credences happens. And if the agent conditionalises over their past experiences of predictions and comes to the conclusion that always the opposite of what they predicted happened, and adjusts their credences to accommodate this, then the world changes again to screw up their predictions once more, and so forth.²⁶ So if nature is irregular, there is no way to tell which way of predicting is the most successful. But what we do know is that regardless of whether the world is regular or not, conditionalising maximises expected epistemic utility.²⁷

But is that enough, is that all we expect from a justification of induction, given that we accept a justification of conditionalisation as a justification of inductive practice? I argue that it is, and that it is all we can, or even *should* hope for. That the world might be (or rather, is) such that conditionalisation sometimes produces false predictions is exactly what we should expect from a defeasible form of inference. That induction might fail is the very nature of induction. We should not fall into the same trap as people who try to demonstrate the truth-conduciveness of induction and deny the special defeasible character of induction: induction can lead from true premises to a false conclusion. We have to allow that by inferring inductively, we might get things wrong; not just occasionally, but often, or even always, if we are very unlucky and live in a particularly nasty world. The important task is to show why we would still be justified to infer inductively, even if we cannot be certain that the conclusion will (likely) be true. The expected epistemic utility justification does just that by demonstrating that conditionalisation maximises expected epistemic utility, even if our predictions might turn out false. So if we want to avoid denying that induction is *sui generis*, and if we want to avoid Hume's dilemma, this is the only possible option we have. Since for this justification, it is irrelevant whether our predictions actually turn out true, there is a way around Hume's dilemma in a way there was none for attempts to justify induction that are designed to show that induction leads towards the truth.

The same goes for the Dutch strategy argument. It, too, does not rely on whether there is any regularity behind what happens in our world, since it only shows that conditionalisation produces coherent credence functions, not that they in any way correspond to any worldly regularities. So as controversial as Dutch strategy arguments are as a tool to justify conditionalisation, they have very few restrictions on the regularity of nature: the agent themselves and the bookie must not change the bets, and the agent must not change their updating norm. Otherwise, Dutch strategies are just measures of whether the agent holds coherent beliefs, and that is independent from whether nature remains constant. Since it is entirely in the agent's hands whether they prefer coherent credence functions, this defence stands.

²⁶This comes very close to Popper's case against epistemic determinism, where he argues that it is impossible to make predictions of a physical system we ourselves are part of, because the prediction itself changes the system, which leads into an infinite regress of factoring in the physical change predictions make in order to make predictions (Popper 1982).

²⁷Whether conditionalisation delivers actual or merely expected epistemic accuracy is discussed, e.g. in Douven (2013).

While these defences of conditionalisation show that a justification of inductive practice which does not rely on the uniformity of nature is possible, neither conditionalisation nor its defences are entirely uncontroversial. There is a sizeable debate whether conditionalisation is actually the best way to update one's credences, even in a Bayesian setting. Bas van Fraassen, for example, proposes two different updating rules, which are both a bit weaker than conditionalisation in the sense that both are implied by conditionalisation, but not vice-versa: special and general reflection.²⁸ Moreover, while arguments such as the expected epistemic utility and the Dutch-book argument may not rely on the uniformity of nature, conditionalisation itself may come with some other metaphysical preconditions. Michael J. Shaffer, for example has argued that conditionalisation, since it involves assertions about an epistemic agent's future credences, requires a view of the future in which the truth value of future contingents does not routinely turn out being false or indeterminate in order for conditionalisation to be coherent.²⁹ While I do not want to get too deep into the discussion of alternative updating rules, two things can be said in reply to this challenge. Firstly, while conditionalisation may, if Shaffer's argument is correct, come with this metaphysical demand about the possibility that future contingents can possibly have a positive truth value, at least the *justification* of conditionalisation does not require a principle such as the uniformity of nature that can only be empirically (or, more specifically, inductively) established, which would be circular. And secondly, I agree that there might be alternative updating rules to conditionalisation, such as van Fraassen's special and general reflection, which each come with their own attempts at a justification.³⁰ I do not hold a stake in this debate. The attempts to justify conditionalisation discussed above are meant to serve as an example of how a justification of induction might be possible, not as an argument that conditionalisation is a better updating rule compared to its rivals. As long as the justification of it can escape Hume's dilemma, any alternative updating rule has a chance of being justifiable, if it does not fail for other reasons.

So if we accept that rationality arguments for updating rules such as conditionalising are justifications of inductive practice, these rationality arguments seem to be the only way forward for a possible justification of induction.

6 Conclusion

The debate about induction has for a long time been dominated by accounts, sceptical and positive alike, which presupposed standards for justification that are impossible to satisfy without falling prey to some version of Hume's dilemma. The recent arguments for conditionalisation can be seen as a way to justify our inductive practice by demonstrating why it is rational to update our beliefs according to new evidence. Crucially, our rationality to do so does not depend on the regularity of nature, and

²⁸See, e.g., van Fraassen (1984, 1995)

²⁹See Shaffer (2014).

³⁰See, e.g. van Fraassen (1984, 1995).

hence evades Hume's dilemma. If we accept arguments for the rationality of inductive practice as a justification of induction, then this is the only way forward for a possible justification. To forego any attempt to show how induction could lead to true conclusions, or how it is likely that the conclusion is true, is no cop-out: it goes against the very nature of induction as a defeasible type of reasoning to even try that. Instead, if we want to justify induction, we would have to demonstrate that we are rational to engage in inductive reasoning. The Dutch strategy and expected epistemic utility arguments do just that. I do not claim that the old riddle has been solved. In order to make such a claim, I would have to argue that the Dutch strategy and the expected epistemic utility arguments actually show what they are designed to show, and I will refrain from judgement regarding that matter. However, *if* we wanted to solve the old riddle by actually proposing a justification of induction, demonstrating the rationality of inductive practice is the only available option.

Acknowledgments I would like to thank the anonymous reviewers and the participants of a number of conferences where I presented this paper for their helpful and constructive criticism. I am especially grateful to Johanna Thoma, Thomas Müller, Ben Eva, Antje Rumberg, Verena Wagner, Martin Rechenauer, Tobias Henschen, and Dominic Alford-Duguid. I am particularly grateful to the Deutsche Forschungsgemeinschaft (DFG) for their generous funding of my project 'The Metaphysics of Induction' [grant number 276459891].

Compliance with Ethical Standards

Conflict of Interest The authors declare that they have no conflict of interest.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- BonJour, L. (1998). *In Defense of Pure Reason. A Rationalist Account of A Priori Justification*. Cambridge Studies in Philosophy. Cambridge: Cambridge University Press.
- Douven, I. (2013). Inference to the best explanation, dutch books, and inaccuracy minimisation. *The Philosophical Quarterly*, 63(252), 428–444.
- Greaves, H., & Wallace, D. (2006). Justifying conditionalization: conditionalization maximizes expected epistemic utility. *Mind*, 115(459), 607–632.
- Haack, S. (1976). The justification of deduction. *Mind*, 85(337), 112–119.
- Hume, D. (2000). In Norton, D.F., & Norton, M.J. (Eds.) *Oxford Philosophical Texts*. Oxford: Oxford University Press. [1739/–40].
- Joyce, J. (1998). A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65, 575–603.
- Leitgeb, H., & Pettigrew, R. (2010a). An objective justification of Bayesianism I measuring inaccuracy. *Philosophy of Science*, 77, 201–235.
- Leitgeb, H., & Pettigrew, R. (2010b). An objective justification of Bayesianism II. The consequences of minimizing inaccuracy. *Philosophy of Science*, 77, 236–272.
- Lewis, D. (1999). *Papers in metaphysics and epistemology*. Cambridge: Cambridge University Press.
- Lipton, P. (2004). *Inference to the best explanation*, 2nd. London: Routledge.
- Karl, R. (1966). *Popper. Die Logik der Forschung*. Tübingen: J. C. B. Mohr (Paul Siebeck).
- Popper, K.R. (1982). *The open universe: an argument for indeterminism. from the postscript to the logic of scientific discovery*. London: Hutchinson.

- Popper, K.R. (2005). *The logic of scientific discovery*. Routledge classics. London: Routledge. [1935] Originally published in English in 1959.
- Briggs, R.A., & Pettigrew, R. (forthcoming). An accuracy-dominance argument for conditionalization. Forthcoming in *Noûs*: <https://doi.org/10.1111/nous.12258>.
- Reichenbach, H. (1935). *Wahrscheinlichkeitslehre – Eine Untersuchung über die logischen und mathematischen Grundlagen der Wahrscheinlichkeitsrechnung*. A.W. Sijthoff's Uitgevermaatschappij N.V., Leiden.
- Salmon, W. (1974). The concept of inductive inference. In Swinburne, R. (Ed.) *The justification of induction* (pp. 49–57). Oxford: Oxford University Press.
- Shaffer, M.J. (2014). Reflection, conditionalization and indeterminacy about the future. *The Reasoner*, 8(6), 65–66.
- Smart, B.T.H. (2013). Is the Humean defeated by induction? *Philosophical Studies*, 162, 319–332.
- Spohn, W. (2012). *The laws of belief: ranking theory and its philosophical applications*. Oxford: Oxford University Press.
- Spohn, W. (2016). Enumerative induction. In Beierle, C., Brewka, G., Timm, M. (Eds.) *Foundations of formal rationality: essays dedicated to Gabriele Kern-Isberner on the occasion of her 60th birthday* (pp. 96–114). London: College Publications.
- Stove, D. (1986). *The rationality of induction*. Oxford: Oxford University Press.
- Strawson, P.F. (1952). *An introduction to logical theory*, Methuen, London.
- Teller, P. (1973). Conditionalization and observation. *Synthese*, 26, 218–258.
- van Fraassen, B. (1984). Belief and the will. *Journal of Philosophy*, 81, 235–256.
- van Fraassen, B. (1989). *Laws and symmetry*. Oxford: Clarendon Press.
- van Fraassen, B. (1995). Belief and the problem of ulysses and the sirens. *Philosophical Studies*, 77, 7–37.
- Vineberg, S., Dutch book arguments, Zalta, E.N. (2016). In *The Stanford Encyclopedia of Philosophy*. *Metaphysics Research Lab, Stanford University spring 2016 edition*.
- Williams, D. (1947). *The ground of induction*. Cambridge: Harvard University Press.