This is the penultimate draft of a paper forthcoming in: *The British Journal for the Philosophy of Science*. https://www.journals.uchicago.edu/doi/abs/10.1086/724449

# In Defence of Science: Two Ways to Rehabilitate Reichenbach's Vindication of Induction

Jochen Briesen

**Abstract**

Confronted with the problem of induction, Hans Reichenbach accepts that we cannot justify that induction is reliable. He tries to solve the problem by proving a weaker proposition: that induction is an optimal method of prediction, because it is guaranteed not to be worse and may be better than any alternative. Regarding the most serious objection to his approach, Reichenbach himself hints at an answer without spelling it out. In this paper, I will argue that there are two workable strategies to rehabilitate Reichenbach's account. The first leads to the widely discussed method of meta-induction, as proposed by Gerhard Schurz. The second strategy has not been suggested thus far. I will develop the second strategy and argue for it being, in some respects, superior to the first and closer to Reichenbach's own position. The strategy is based on Reichenbach's idea that the inductive straight rule is not only applicable on the object but also on the method level. He does not spell out how exactly this insight is supposed to save his account. But he seems to assume that nothing more than the straight rule and the different levels of its application is needed for this purpose. The strategy introduced in this paper illustrates that this assumption is correct.

## 1 Introduction

Inductive reasoning is a central tool in science. The following are simple schemata of certain forms of inductive inferences (enumerative induction (I), (II); statistical induction (III), (IV)):

| I) | II) | III) | IV) |
|---|---|---|---|
| (Observed) $X_1$ is $F$<br>(Observed) $X_2$ is $F$<br>(Observed) $X_3$ is $F$<br>…<br>(Observed) $X_n$ is $F$<br><br>Thus: (Unobserved) $X_{n+1}$ is $F$ | (Observed) $X_1$ is $F$<br>(Observed) $X_2$ is $F$<br>(Observed) $X_3$ is $F$<br>…<br>(Observed) $X_n$ is $F$<br><br>Thus: All $X$ are $F$ | $p$% of the observed $X$ are $F$<br><br>Thus: With probability $p$, the next unobserved $X$ is $F$. | $p$% of the observed $X$ are $F$<br><br>Thus: $p$% of all $X$ are $F$. |

As illustrated in these schemata, inductive inferences draw conclusions from the previously observed to the yet unobserved; in this sense induction is a method of prediction. Some forms of inductive inference predict something about the next previously unobserved instance of an event (see (I), (III)), while others predict something about all of them (see (II), (IV)).

However, whether induction is a *reliable* method of prediction—that is, whether it leads from true premises to true conclusions most of the time—depends on the uniformity principle. The uniformity principle states (approximately) that the hitherto-unobserved cases of a phenomenon resemble those cases hitherto observed. It can be argued, however, that this principle is justified neither a priori nor a posteriori. It is not justified a priori because to justify that the empirical world behaves uniformly, we have to look at it. There are simply no a priori reasons to assume the uniformity of nature. Furthermore, the principle is not justified a posteriori because any attempt to do so would be circular (or would lead into a regress): such an attempt would point to the fact that we have experienced nature behaving uniformly in the past and predict that it will continue to do so in the future—thereby presupposing that which is supposed to be justified in the first place, namely that the world behaves uniformly. However, because every justification is either a priori or a posteriori, the uniformity principle and thereby the reliability of induction is not justified at all. This is the problem of induction.[1] In times of growing scepticism about science, this fundamental problem in epistemology also gains relevance from a social and political point of view. A satisfactory solution to the problem should therefore be dialectically effective in that it should be able to convince sceptics with respect to science and adherents of other nonscientific prediction methods.

How can we solve the problem of induction? Is there a way to block the argument that leads to the problem? Throughout the history of philosophy, various affirmative answers to this question have been suggested. Immanuel Kant ([1998]), for example, famously argued that there are indeed a priori reasons to assume that the experientially accessible world behaves uniformly. Others have defended an a posteriori justification of induction by arguing that the kind of circularity involved (namely, rule-circularity in contrast to premise-circularity) is in fact epistemically acceptable (Papineau [1992]; Van Cleve [1984]). Still others have accepted that Hume's argument cannot be blocked but take the position that science ultimately does not depend on inductive procedures (Popper [2002]). None of these widely discussed reactions, however, has been met with widespread approval.

In this paper, I will focus on the vindication of induction suggested by Hans Reichenbach ([2006], §31; [1949], §91). Reichenbach accepts the conclusion of Hume's sceptical argument; that is, he accepts that we cannot justify that induction is a reliable method of prediction. However, in his view we can justify that induction is optimal in the following sense: If there is any method that reliably infers from the observed to the unobserved, it is the inductive method. This attempted defence of induction does not amount to a justification of the claim that induction is reliable, but to something weaker: induction is in any case not worse, but possibly better than any alternative procedure—for example, believing self-proclaimed clairvoyants, faith healers, religious leaders, or political demagogues.

Even if Reichenbach's proposal leads to something weaker than we might have hoped for, the proposal, if successful, would still amount to an important achievement. This is especially true in view of the increasing non-academic relevance of the problem. If Reichenbach is right, then we can rationally defend induction against attacks from advocates of nonscientific methods. If he is correct, then with respect to inductive

[1] For the classical source of the problem, see (Hume [2009], Book 1, part iii, sec. 6; [2010], sect. iv); for a useful overview to the discussion concerning the problem, see (Henderson [2018]).

methods—and only with respect to them—it can be proven that they are guaranteed to be as good as alternative methods, and possibly better.[2]

Unfortunately, Reichenbach's proposal has a number of difficulties, many of which he was aware of. With respect to what is probably the most serious problem, he also hinted at a solution. However, he never spelled out a solution in detail. In this paper, I will argue that there are basically two strategies to work out Reichenbach's vague remarks on how to rehabilitate his initial solution to the problem of induction. The first strategy leads to the method of meta-induction proposed by Gerhard Schurz, which has received much attention—and rightly so. Schurz's reflections and the results concerning meta-induction have far-reaching and fascinating consequences for a number of epistemological issues (Schurz [2008]; [2018]; [2019]; [forthcoming]). In contrast, the second strategy has not been discussed thus far. In this paper, I will develop this second strategy and argue for it being, in some respects, superior to the first.

To avoid misunderstandings, it should be mentioned at the outset that I cannot discuss—let alone solve—every problem that arises with respect to inductive reasoning. For example, I will not address the so-called 'new-riddle of induction' raised by Goodman ([1955], sect. 3.4).[3] My focus is solely on addressing the main difficulties of Reichenbach's interesting response with respect to the aforementioned induction-sceptic argument.

In Section 2, I will first outline Reichenbach's solution and address the central difficulties of his proposal. Additionally, I will present his own hints on how to solve the most serious problem of his approach. In Section 3, I will introduce the two strategies of spelling out Reichenbach's hints. In section 4, I will discuss the first strategy in more detail and, in this context, introduce the main features of Gerhard Schurz's meta-induction. Then, in Section 5, I will discuss the second strategy and establish its advantages over the first. I will end the investigation with a brief summary in Section 6.

## 2 Reichenbach's Vindication of Induction

Reichenbach accepts that there are neither a priori nor a posteriori reasons for believing that induction leads reliably from true premises to true conclusions. However, he argues that it is nevertheless epistemically rational to use the method of induction. He argues that with respect to our epistemic goal of maximizing the set of true beliefs, induction is still the best method available to us. It is optimal in the following sense: With respect to our epistemic goal, it is guaranteed that the method is at least not worse and possibly better than alternative procedures. If the world behaves uniformly then induction is better, and if the world does not behave uniformly then induction is at least not worse than alternative noninductive methods. Thus, within Reichenbach's suggestion, it is not the belief in the reliability of induction that is justified, but the decision to hold on to induction as a particular belief-forming method. That is why Reichenbach's defence is sometimes called a 'pragmatic vindication of induction' (Feigl [1950]). It is important to note, however, that Reichenbach's best alternative approach remains an epistemic defence insofar as the

[2] Note that Reichenbach does not claim that it can be proved that inductive methods are the only methods that are optimal in this respect (i.e., he does not argue that induction is a dominant prediction strategy). He rather claims that it is the only method for which it can be proven that it is optimal (Schurz [2008], p. 303).
[3] For interesting recent discussions of this problem see, for example, (Freitag ([2015]; Zinke ([2020]).

inductive method is supposed to be optimal with respect to our epistemic goal of maximizing the set of true beliefs and not with respect to any other non-epistemic aim.

Why does Reichenbach think that induction is optimal? The answer to this question depends on his characterization of uniformity and his specification of the inductive method, both of which in turn depend on his probabilistic epistemology together with a frequentist interpretation of probability.

Suppose we have observed that out of 1,000 people who ate a certain kind of mushroom—let's say a death cap—800 people died. Thus, we have observed the relative frequency $f_n$ of an incident *m* (dying-after-death-cap-consumption) in *n* cases (of death-cap-consumption). In our example: $f_{1000}$=8/10=0.8. Reichenbach introduces the notion of a 'limiting relative frequency' as follows: 'The frequency $f_n$ has a limit at *p*, if for any given sufficiently small number $\epsilon$ there is an *n* such that $f_n$ is within p $\pm\, \epsilon$ and remains within this interval for all the rest of the series' (Reichenbach [1938/2006], p. 351). Thus, $f_n$ has a limit, if for any extension of the series of observations, at some point the frequency will continue to fall within a small interval.

If the world behaves uniformly, then such limiting relative frequencies exist. Let us suppose that we have reached the limit in our example after 1,000 observed incidents, then the limiting relative frequency of dying-after-death-cap-consumption is close to 0.8; in Reichenbach's words: it is within the interval 0.8 $\pm\, \epsilon$. The frequency will not change drastically any further, no matter how many more observations we will make in the future. If, on the other hand, the world does not behave uniformly, then no limiting relative frequencies exist. In this case the relative frequency of incident *m* (dying-after-death-cap-consumption) in *n* cases (of death-cap-consumption) can change drastically with any extension of the series (Reichenbach [2006], p. 350).

Thus, Reichenbach specifies what it means to say that the world behaves uniformly via recourse to the notion of 'limiting relative frequency'. To understand his vindication of induction, we also have to clarify how he characterizes the method of induction. In his view induction is 'a procedure in which the relative frequency observed statistically is assumed to hold approximately for any future prolongation of the series' (Reichenbach [2006], p. 340). This method simply demands that we take the current observed relative frequency as the limiting frequency. Reichenbach calls this the 'the principle of induction', but it is today best known as the 'straight rule'. In the context of a frequentist interpretation of probability, this characterization fits well with the schemata of inductive inferences given in section 1—where enumerative inductive inferences (I)–(II) can be understood as special cases of statistical inductive inferences (III)–(IV).

Of course, the straight rule will deliver false results at the beginning of the series of observations, because it will map random fluctuations in the sample frequency. In the long run, however, the straight rule is guaranteed to determine the correct limiting relative frequency—provided there is such a limit. With Reichenbach, we can therefore hold that if the world behaves uniformly—that is, if there are limiting relative frequencies, then the inductive method is reliable—that is, then the straight rule will correctly determine the limit in the long run. Furthermore, if the world does not behave uniformly, then no limiting frequencies exist, and no method will find them.

We are now in a position to formulate the argument that is supposed to establish the optimality of induction:

(i) Either the world behaves uniformly and there are, thus, limiting relative frequencies, or not.

(ii) If the world behaves uniformly and there are limiting relative frequencies, then induction is in the long run a successful method of prediction (i.e., then induction is guaranteed to determine the limiting relative frequencies in the long run).

(iii) If the world *does not* behave uniformly and there are *no* limiting relative frequencies, then neither inductive nor noninductive methods of prediction are successful.

(iv) Therefore, induction is the optimal method of prediction: it is guaranteed not to be worse and possibly better than noninductive methods.

As a defence of induction, this argument invites several objections. First, in the argument the goal of inductive inferences is restricted to finding limiting relative frequencies. However, not all forms of scientific reasoning can be reduced to this task. Thus, the argument cannot prove that scientific reasoning in general is optimal. This remark does not amount to a devastating objection. Reichenbach's argument is not strong enough to establish optimality for all kinds of scientific inferences, but it might still be strong enough to establish optimality for the kinds of induction introduced in section 1—after all, they can be understood as being directed at estimating limiting frequencies. Because these kinds of induction are undeniably important to science, their defence would remain an important achievement.

Second, Reichenbach's defence seems too dependent on the long run. He can only show that the inductive method (straight rule) will determine the correct limit—if there is one—in the long run. However, after any number of observations, no matter how large, it is always possible that we will not have reached the limit. Even worse, it is always possible that the observed frequency is maximally far away from the true limit. Thus, we never really know what the limit is (BonJour [1998], p. 194; Salmon [1966], p. 53; Schurz [2019], p. 82; Skyrms [1964], pp. 259–60). This observation is correct and it would certainly be nice if the dependence on the long run could be avoided somehow. Nonetheless, I do not consider this a devastating objection. After all, even in the face of this objection it remains true: if there is a limit, the method is guaranteed to find it (in the long run). Thus, using the method is still rational and justified from an epistemic perspective. Pointing to the fact that we never know whether we have reached the limit is less an indication of the nonrationality of the inductive method than an indication of the open-endedness of science. (Reichenbach's ([2006], pp. 361–62; [1949], pp. 447–48) own response to the problem consists in introducing a 'practical limit').

Third, the achieved defence does not only concern the inductive method—that is, the straight rule—but many other methods as well, namely all methods that converge the straight rule asymptotically. For simple examples of such asymptotic methods, add any function to the straight rule that converges to 0 with increasing $n$ (Henderson [2018], sec. 7.1; Salmon [1966], p. 53). This is a serious problem. How can we establish that the straight rule is in no case worse than one of the asymptotic rules? Isn't it possible that one of the asymptotic rules finds the limiting frequency faster or more accurately (with smaller $\epsilon$) than the straight rule? In my view, even in the face of this concern, using the inductive method remains justified, because it is still guaranteed that, if a limiting frequency exists, then the straight rule will find it (in the long run). Furthermore, the other methods are more complex, and they only work because they gradually approach the inductive method. Thus, the inductive method (straight rule) seems more fundamental and descriptively simpler than the other methods, and, therefore, with regard to the epistemic aims of simplicity and explanatory strength as well as with regard to its application, even superior to them (see also Reichenbach [1949], pp. 475–76). (Reichenbach ([2006], pp. 355–56) additionally argues that the inductive straight rule is less risky.)

Fourth, Reichenbach's argument depends on an overly simplistic characterization of our epistemic goal as well as an incomplete description of the available options. The goal is determined as maximization of true beliefs, and the options are limited to the choice between inductive or noninductive methods of belief formation. This is misleading. Our epistemic goal is not simply to maximize the number of true beliefs, otherwise it would be epistemically rational to believe everything whatsoever. After all, this would guarantee the maximum amount of true beliefs. To avoid this implausible consequence, our epistemic goal must be characterized in more complex terms: increasing the number of true beliefs while at the same time avoiding false ones. This more complex goal, however, brings into play the hitherto ignored option of suspending judgment. If it is true that in case of the nonuniformity of nature neither inductive nor noninductive methods are successful, then in this case, with respect to our more complex epistemic goal, suspension of judgment would be better than following the inductive method. After all, in this case no option would increase the set of true beliefs, but suspension of judgment would be the only option that would at least not increase the set of false ones. Thus, within an accurate description of our epistemic aim and the available options, the inductive method cannot be proven to be optimal anymore (for a related concern, see Lange [2011], p. 77).

This also is a serious concern with respect to Reichenbach's suggestion. Probably the only way to deal with it is by referring to the practical necessity of making predictions. Given the fact that to survive we have to make predictions and infer from the observed to the unobserved, it is reasonable to ignore the option of suspension. Often suspension is simply not a viable option for us. We have to accept, though, that ignoring the option of suspension cannot be motivated from a strictly epistemic point of view, that is, in terms of our epistemic goal alone. Thus, calling Reichenbach's position a 'pragmatic' vindication of induction seems appropriate in another sense as well: on closer examination, his epistemic defence of induction depends on the assumption of certain practical necessities.

The fifth and final objection is the most devastating one: premise (iii) is simply false. It is not true that no method of prediction, neither inductive nor noninductive, can be successful if the world does not behave uniformly (Herz [1936]; Skyrms [1964], p. 260). Even in a world in which no limiting relative frequencies exist, it could still be true that soothsayers with perfect foresight reliably predict events in the future. Think of the mushroom example again. If the world is so disorderly that there is no limiting relative frequency of dying after death-cap consumption, then a soothsayer could not determine the limiting relative frequency—after all, it does not exist—but she could still successfully and reliably predict whether the next person who eats a death cap dies or not. Not all prediction methods are tied to finding limiting relative frequencies.

Reichenbach was keenly aware of this problem. He replied that a world with reliable soothsayers would exhibit a certain form of uniformity again, which could be specified by a limiting relative frequency and which, therefore, inductivists could make use of. If the soothsayer is reliable in making predictions, then the limiting relative frequency of correct predictions is above a certain threshold and the straight rule is guaranteed to find it (Reichenbach [2006], pp. 358–60; [1949]; p. 476). This is true. On the level of prediction methods, the inductive straight rule is guaranteed to determine the relative frequency of correct answers within a prediction method in the long run—provided there exists such a limit.

All this said, how is this supposed to help us vindicate induction on the object level: that is, induction applied not at the level of methods, but at the level of events

(Schurz [2008], p. 281; [2019], pp. 82–3; Skyrms [1975], p. 44)? This is a crucial question because our goal, of course, was to prove that it is epistemically rational to use the scientific method of induction on the object level. It is therefore all the more surprising that Reichenbach himself does not even address this question.

## 3 Two Ways to Rehabilitate Reichenbach's Solution

Induction can be used on the object level or on the method level. Induction at the object level amounts to using induction to make predictions of events in the world. In our mushroom example, the inductive straight rule is used at the object level to predict the probability of death-after-death-cap-consumption. Induction can also be used at the method level. In our soothsayer example, the inductive straight rule is used at the method level to predict the limiting relative frequency of correct answers of the soothsayer prediction method. I will refer to induction on the method level as 'method-induction'. The central question is: How can Reichenbach's vague hints concerning induction on the method level be spelled out as relevant to induction on the object level? Two strategies can be distinguished.

(I) Inductivists must take into account the success of alternative noninductive methods and base their own predictions on the success of those other methods. This new and more complex inductive method is called 'meta-induction' (Schurz [2008]; [2019]) and must be specified in detail. It has little in common with the straight rule favoured by Reichenbach; nevertheless, it can be shown that this more complex inductive method is optimal in the previously specified sense.

(II) Inductivists apply the straight rule not only at the object level, but also at the method level. In a first step, it can be shown a priori that the inductive straight rule is optimal with respect to the method level. In any world—uniform or not—the inductive straight rule is the best option to determine which method on the object level is reliable—given there is a reliable method in the first place. In a second step, inductivists apply the straight rule to the prediction methods actually used in our world. This application leads to the result that the scientific inductive methods are more reliable on the object level than noninductive alternatives. This second step can be characterized as an inductive, a posteriori justification of object induction. In contrast to former attempts to justify induction a posteriori, the suggested strategy, however, is neither circular, nor does it lead into a regress. After all, in the first step induction at the method level is proven to be optimal a priori.

Strategy (I) has been explored by Gerhard Schurz. As already indicated, he calls the more complex inductive method 'meta-induction'. Regardless of the name, within the framework of meta-induction predictions are made with respect to events on the object level. The difference from the usual forms of induction is that meta-induction takes into account the predictions and success rates of all other prediction methods. The usual inductive methods that predict events without basing their predictions on the predictions of all other methods are called 'object induction'. Thus, in the context of discussing strategies (I) and (II), we should differentiate among (A) object induction, which predicts events on the object level, without basing predictions on the predictions and success rates of other methods; (B) meta-induction, which predicts events on the object level by basing their predictions on the predictions and success rates of all accessible prediction methods; and (C) method induction, which does not make predictions of events on the object level but predictions concerning the reliability of other prediction methods.

Referring to mathematical results in the field of computational learning theory, Schurz proves that there are meta-inductive methods that are optimal. This is an impressive result. However, the meta-inductive method of prediction is greatly different from the usual inductive methods used in science.[4] Therefore, in the context of strategy (I), Schurz also relies on a second step. He argues: if we apply the optimal method of meta-induction in our world, then we see that the actual scientific inductive procedures are reliable—after all, meta-induction keeps track of the success rates of all methods on the object level.

However, if strategy (I) also relies on a double step in defending the inductive inferences used in science, it seems worthwhile to examine strategy (II) more closely. In the following, I will first outline Schurz's version of strategy (I) in a little more detail (see section 4). I will then investigate strategy (II) and argue that it is in certain respects superior to Schurz's account (see section 5).

## 4 Strategy (I): Schurz's Meta-Induction

Meta-inductive methods MI predict events by taking into account the predictions of all other methods. A simple example is the following version of an imitate the best meta-induction (ITB*):[5]

> ITB*: Observe which method (or methods) predicted correctly in the last round and follow it (or one of it) in the next round of predictions.

ITB* is not optimal. There are worlds in which ITB* is worse than other methods. Consider two methods, M1 and M2, that always make opposing predictions. Furthermore, assume the world behaves in such a way that whenever M1 is correct (and M2 is wrong), in the next round of predictions M2 is correct (and M1 is wrong), and so on. If the methods M1 and M2 alternate in this way, both are correct 50% of the time. However, a person who follows ITB* will make wrong predictions 100% of the time. She will observe that M1 is right in a given round of predictions, she will follow M1 in the next round, and she will predict wrongly. Furthermore, she will have observed that in this second round M2 was correct, so she will follow M2 in the third round, and will be wrong again, and so forth. Therefore, ITB* is not an optimal method of prediction; it is not guaranteed for any course of events that ITB* is not worse than alternative methods.

---

4 This is not to say that scientists never use methods that aggregate the results of other methods. Consider, for example, certain forms of meta-analysis in statistics.

5 Note that Schurz characterizes imitate the best meta-induction (ITB) differently. This is why I do not use 'ITB', but 'ITB*'. Furthermore, Schurz ([2019], sec. 6.3) argues for the more general conclusion that not only ITB, but every one-favourite prediction method is non-optimal.

The central idea of an optimal meta-inductive method lies in two modifications: first, an optimal meta-induction should not favour a single method but should rather make its predictions depend on the predictions of all (accessible) methods. Second, it should base its predictions on the predictions of all other methods in a special way: to be optimal, its prediction must be the weighted average of the predictions of all the other methods. The framework in which Schurz specifies these improvements is the framework of sequential predictions in computational learning theory. In the following, I will mostly follow Schurz's notation.

Schurz defines a prediction game G as a pair ((e), Π) of a sequence of events (e) and a pool Π of finitely many prediction methods or players: {M1, M2, M3,…}.[6] Events are identified with values in a set of values Val, and predictions are elements in a set of values $Val_{pred}$. For real-valued prediction games, both events and predictions are in the range of [0,1], that is, $Val = Val_{pred} = [0,1]$.

Players have the task of predicting events, where '$pred_n(M)$' denotes a prediction of the method M for time n that is issued at n-1. The loss that a method M incurs in round n is measured by a loss function:

$$loss_n(M) =_{def} loss(pred_n(M), e_n).$$

This function measures the deviation of prediction $pred_n$ from the event $e_n$. Various loss functions are possible, but for the purposes of this paper we will concentrate on the so-called natural loss function:

$$loss(pred_n, e_n) =_{def} |pred_n - e_n|.$$

Via recourse to the loss of M, Schurz defines the 'score' that M earns for $pred_n$ as follows:

$$score_n(M) =_{def} 1 - loss_n(M).$$

Via recourse to the score, the 'absolute success' achieved by M until n is defined as the sum of M's scores for predictions until n:

$$abs_n(M) =_{def} \sum_{i=1}^{n} score_i(M).$$

And via recourse to the absolute success, the 'success rate' of M at n is defined by:

$$suc_n(M) =_{def} \frac{abs_n(M)}{n}.$$

Given these notions, we can finally formulate what we are looking for in precise terms. What we are looking for is a meta-inductive strategy MI that has access to the predictions of all methods M in Π (including its own) and predicts in such a way that is optimal with respect to Π. MI is optimal with respect to Π if it will always, irrespective of the history of events, be at least as successful as any M ∈ Π in the long run. This long run optimality requires that the difference of MI's success rate—its average success per round—and that of any other M converges to a nonnegative value:

[6] The fact that Schurz's meta-inductive strategy is restricted to a finite set of prediction methods has been objected to by Arnold ([2010]). Schurz ([2018], pp. 3891–92) responds convincingly to the objection. For a helpful discussion of this issue, see Sterkenburg ([2019], pp. 986–89).

$$\lim_{n\to\infty}(\mathrm{suc_n(MI)} - \mathrm{suc_n(M)}) \geq 0.$$

There are other and more complex ways to specify the optimality of MI. Not all of them require that the difference between MI's success rate and that of any other M converges to a limit, and not all of them are restricted to long run optimality (Schurz [2008]; [2019]; Sterkenburg [2020]). For now, however, it is sufficient to concentrate on the less complex long run optimality as it is heretofore specified.[7]

Which meta-inductive Method MI achieves the characterized goal of long run optimality? It turns out that an optimal meta-inductive strategy predicts in each round by a weighted average of the predictions of all the other methods. Given the formula for weighted means, the weighted-average meta-inductivist wMI predicts as follows:

$$\mathrm{pred_{n+1}(wMI)} =_{\mathrm{def}} \frac{\sum_{\mathrm{M}\in\Pi} \mathrm{w_n(M)} \cdot \mathrm{pred_{n+1}(M)}}{\sum_{\mathrm{M}\in\Pi} \mathrm{w_n(M)}}.$$

Which weight $\mathrm{w_n(M)}$ do we have to choose so that this prediction method turns out to be optimal? Results in computational learning theory established that the weights have to depend on the strategies' losses. The idea is to assign weights that are determined by how much better M did than the meta-inductive method in hindsight. Schurz calls this M's 'attractivity' and defines it as follows:

$$\mathrm{at_n(M)} =_{\mathrm{def}} \mathrm{suc_n(M)} - \mathrm{suc_n(MI)}.$$

In the method of attractivity-weighted meta-induction awMI, the weights for M are directly based on M's attractivity—provided that M does not have negative attractivity, in which case it is simply ignored:

For all $n \geq 1$ and M, the weight $\mathrm{w_n(M)} = \mathrm{at_n(M)}$ if $\mathrm{at_n} \geq 0$
$= 0$, otherwise.

Thus, the attractivity-weighted meta-inductive strategy awMI predicts as follows:

$$\mathrm{pred_{n+1}(awMI)} =_{\mathrm{def}} \frac{\sum_{\mathrm{M}\in\Pi} \mathrm{at_n(M)} \cdot \mathrm{pred_{n+1}(M)}}{\sum_{\mathrm{M}\in\Pi} \mathrm{at_n(M)}}.$$

This method turns out to be optimal in the previously specified sense. One can prove that it satisfies the specified condition for an optimal strategy (see Cesa-Bianchi and Lugosi [2006]; Schurz [2008]; [2019]).

It is important to note that comparable results can be established for various prediction games (not only with real-valued but also with binary (discrete) events) and different notions of optimality (long-run and short-run optimality). Demonstrating that awMI is not only optimal in the long run but also optimal in the short run is particularly important. This result allows Schurz to avoid two problems that Reichenbach's original proposal faces. First, it avoids the problem of being too dependent on the long run (see the second problem in section 2). Second, it avoids the problem of asymptotic methods (see the third problem in section 2). There are many meta-inductive methods that converge

[7] I will discuss the results with regard to short-run optimality, which are of central importance for Schurz in due course (see Schurz [2019], sec. 6.6).

awMI and are, therefore, also optimal in the long run (while differing drastically from awMI in their short-run predictions). However, there is only one variant of meta-induction that is optimal in the short run, namely awMI. Thus, the short-run results are also relevant in sidestepping the problem of asymptotic methods.

In most general terms we can conclude: the meta-inductive method that predicts in a way that at each point in time, it favours the predictions of all other accessible methods to the extent of their relative success so far, is an optimal strategy (Schurz [2019]; Sterkenburg [2020]). This is an excellent result, showing a meta-inductive strategy that is optimal and guaranteed to be no worse than any other method of prediction. What, however, follows from this result for the inductive methods used in science? The meta-inductive strategy awMI is different from many inductive methods used in science. Even though there are various forms of meta-analysis in the natural sciences, usually the inductive methods used in science are not determined by the predictions of all other methods but rather by the history of events on the object level. Thus, many inductive methods used in science are forms of object induction OI. How is the optimality of awMI supposed to justify OI? Schurz acknowledges that a defence of OI requires a second step.

This second step can be summarized as follows: irrespective of the history of events—that is, irrespective of whether the world is uniform or not—awMI is optimal. Thus, irrespective of the world we live in, using awMI is epistemically rational. This is a mathematical result that has been justified a priori. However, if for every world it is epistemically rational to use awMI, it is also epistemically rational to use it in our world. By applying awMI to the prediction methods used in our world, we have to keep track of the success rate of all methods in Π and thereby find that OI is more reliable than other noninductive methods.

Of course, whether this actually is the case is ultimately an empirical question, but it seems reasonable to suppose that it is correct. Most of us will agree that science has been quite successful in the past, more so than alternative methods (such as following religious leaders, political demagogues, or self-proclaimed clairvoyants). However, because this is an empirical question, the justification of OI provided by this second step is a posteriori. What is important to realize, however, is that this new kind of a posteriori justification of induction is neither circular nor does it lead into a regress. It is not circular because OI is not justified by applying OI itself, but rather by applying the meta-inductive method awMI. Furthermore, it does not lead into a regress because the potential regress is stopped by justifying awMI a priori via mathematical reasoning. Thus, the familiar problems of a posteriori attempts to justify induction do not apply to the second justificatory step suggested by Schurz.

As interesting and promising as Schurz's line of thought is, as a defence of induction it faces various difficulties: inductive methods used in science have to be modelled as elements in pool Π of a prediction game. This is problematic for at least two reasons: first, to be an element of Π, the predictions have to be modelled as values in the range of [0,1]. A natural way to do that is to interpret scientific predictions as probabilistic predictions (Schurz [2019], ch. 7). Although this is acceptable to Bayesians, it may not appeal to others (Sterkenburg [2020], p. 526).[8]

[8] One way to circumvent this problem would be to determine optimality not for a single meta-inductive method, but for collective of meta-inductivists (Schurz [2008], pp. 297–99; Schurz [2019], ch. 6.7.2). Schurz ([2019], p. 148) also generalizes his account to discrete predictions without numerical structure in the form of several theorems. I cannot discuss these theorems here in detail.

Second, inductive methods in sciences are not restricted to the aim of making correct predictions of particular events; many inductive methods in science are directed at general, law-like statements. However, inductive methods that aim at law-like generalizations cannot be vindicated by Schurz's suggestion. Meta-induction can justify the predictions of object-induction for the respective next round, but it cannot justify law-like generalizations (Arnold [2010], p. 591; Sterkenburg [forthcoming], p. 6).

Third, in the context of Schurz's proposal we seem to lose the motivation to justify object induction. After all, Schurz has convincingly argued that there is a method that is optimal, namely awMI. So why should we even bother to justify object induction OI used in the sciences by applying awMI? In view of Schurz's results, would it not be rational to switch to awMI in the sciences as well?

Fourth, how exactly is OI justified by applying awMI? The meta-inductivist will give the highest weight to the predictions of OI. If OI is indeed much more successful than all the alternatives, the predictions of awMI will (approximately) coincide with OI's predictions. So what can be justified at most by applying awMI is following OI's current predictions. However, it is at least unclear whether justifying current predictions of a method is the same as justifying the method of object induction OI (Sterkenburg [2020], p. 538).

Fifth, Schurz's defence of scientific induction depends on mathematical considerations in a particular field of computational learning theory (predictions with expert advice). As a defence of induction against attacks from proponents of other methods, it faces the problem of being too complex. In times of emerging scepticism toward science in parts of our society, the defence may be too complex to be dialectically effective.[9]

Although the first two difficulties in particular seem fairly serious to me, the problems listed do not make Schurz's proposal worthless. Possibly some of the problems can be solved, and possibly others must be accepted as necessary theoretical costs for the greater epistemic good. In view of the difficulties, however, I want to suggest that it is worthwhile to take a closer look at strategy (II) with regard to a rehabilitation of Reichenbach's ideas. This is especially true in light of the fact that ultimately even Schurz's proposal of strategy (I) depends on a double argumentative step that combines a priori and a posteriori considerations.

## 5 Strategy (II): Sticking to the Straight Rule in Method Induction

The predictions of method induction do not concern events on the object level but the reliability of prediction methods. I will reserve 'M' for methods that predict events on the object level. What does it mean to say that M is reliable? M is reliable if and only if the limiting relative frequency of its correct predictions relative to all its predictions is above a certain threshold. In the context of strategy (II), it does not matter where exactly the threshold is. Let us assume that it is above 0.8. Thus, M is reliable if and only if the limiting

[9] Sterkenburg ([2019]) discusses another potential problem for Schurz's approach, namely that he cannot show the optimality of meta-induction for an expanding pool of prediction methods. The presentation and discussion of this problem would take up too much space, so I do not include it here. I will shortly come back to it, however, in fn. 12. For Schurz's answer to this problem, see (Schurz [2019], sec. 7.3].

relative frequency of its correct predictions is > 0.8. Therefore, if there is a reliable method M, then there exists a limiting relative frequency of correct predictions that is > 0.8.[10]

Furthermore, Reichenbach specifies induction as the 'procedure in which the relative frequency observed statistically is assumed to hold approximately for any future prolongation of the series' (Reichenbach [2006], p. 306). This amounts to the so-called straight rule that simply demands that we take the current observed relative frequency as the limiting relative frequency. What happens if we apply the straight rule to determine whether there is any reliable method M on the object level? Answer: In the long run, the straight rule is guaranteed to determine the limiting relative frequencies of true predictions of M—provided there is such a limit. Thus, if there is a reliable method M—that is, if there is a limiting relative frequency of correct predictions that is > 0.8—then applying the inductive straight rule will correctly determine the limit in the long run. Furthermore, if there is no reliable method M on the object level, then no procedure on the method level will be able to find a reliable method on the object level. Both conditionals are justified a priori because they directly follow from the formulation of the straight rule together with the specification of the reliability of a prediction method.

In what follows, the term 'method induction' denotes an application of the straight rule to find the reliable methods on the object level. Relying on an argument that is structurally analogous to Reichenbach's argument (i)–(iv) (see section 2), we can prove that method induction is optimal with respect to the goal of finding reliable methods on the object level.

(i)* Either there are reliable methods on the object level and, thus, limiting relative frequencies of correct predictions that are > 0.8 exist, or not.

(ii)* If there is a reliable method M, and, thus, a limiting relative frequency of correct predictions that is > 0.8 exists, then method induction will determine the limit in the long run and will thereby successfully identify the reliable method M.

(iii)* If there is no reliable method M and, thus, no limiting relative frequency of correct predictions that is > 0.8 exists, then neither method induction nor noninductive procedures will find a reliable method M.

(iv)* Therefore, method induction is optimal: it is guaranteed not to be worse and possibly better than noninductive procedures with respect to the goal of finding reliable methods on the object level (in the long run).

---

[10] Suppose there is a method M* which after a certain point achieves a relative frequency of true answers that consistently remains above 0.8 but does not have a stable limit because the frequencies continue to oscillate in the range of [0.81, 0.99]. Isn't M* reliable? And if so, isn't this a problem for the suggested line of thought, because in this case no limiting relative frequency exists? I see two options to deal with this case. (a) Biting the bullet: Since M* does not converge to a limit, there is no limiting relative frequency and, thus, according to the paper's definition of 'reliability', M* is not reliable. (b) Weakening: We have introduced the notion of a limiting relative frequency with Reichenbach as follows: The frequency $f_n$ has a limit at $p$, if for any given sufficiently small number $\epsilon$ there is an $n$ such that $f_n$ is within $p \pm \epsilon$ and remains within this interval for all the rest of the series. However, as far as I can see the basic idea of the strategy would also work with a weaker notion, let's call the weaker notion 'limit*'. According to such a weaker notion, $\epsilon$ would not be required to be chosen arbitrarily small: The frequency $f_n$ has a limit* at $p$, if there is a sufficiently small number $\epsilon$ (say $\epsilon \leq 0.09$) such that there is an $n$ such that $f_n$ is within $p \pm \epsilon$ and remains within this interval for all the rest of the series. According to such a weaker notion, the oscillating method M* has a limit* at $p = 0.9$ with $\epsilon = 0.09$. Via recourse to limit* we could define reliability* (as well as the straight rule*) and reformulate the strategy with these weaker notions.

We have already discussed the problems of Reichenbach's original argument (i)-(iv) (see section 2). The most devastating objection concerned the following premise:

(iii) If the world does not behave uniformly and there are, thus, no limiting relative frequencies (on the object level), then neither inductive nor noninductive methods of prediction are successful.

Premise (iii) is false because not all possible prediction methods on the object level are tied to finding limiting relative frequencies. Even if no limiting relative frequencies on the object level exist, a soothsayer with perfect foresight could still successfully predict upcoming events.

Premise (iii)* does not face this problem. All possible procedures on the method level—that is, all procedures that make predictions concerning the reliability of a method M—are necessarily tied to finding limiting relative frequencies. This is owing to the given specification of 'reliability'. Thus, in contrast to (iii), (iii)* is correct: if there are no limiting relative frequencies of correct predictions on the object level, then no method on the object level is reliable, and no procedure on the method level is able to find a reliable method on the object level.

Argument (i)*–(iv)* is valid, and the premises (i)*–(iii)* are justified a priori. Thus, via recourse to (i)*–(iv)*, the optimality of method induction is justified a priori. Irrespective of the world we live in (whether it contains reliable methods on the object level or not), applying the procedure of method induction is rational; it is not worse and possibly better than any other procedure on the method level. This concludes the first step of strategy (II).

The second step is analogous to the second step Schurz is forced to take in the context of his meta-inductive strategy (I). It can be summarized as follows: irrespective of the world we live in (whether there are reliable methods on the object level or not), using method induction is epistemically rational. However, if for every world it is epistemically rational to use method induction, it is also epistemically rational to use it in our world. By applying method induction (i.e., the inductive straight rule on the method level) in our world, we find that object induction and scientific methods in general are more reliable than other nonscientific methods—such as believing self-proclaimed clairvoyants, miracle healers, and so on. Just as with respect to the second step in Schurz's argument, whether this actually is the case is an empirical question. However, it seems highly reasonable to suppose that it is correct. However, given this is ultimately an empirical question, the justification offered by this second step is a posteriori.[11]

It is important to note that the suggested a posteriori justification is no more circular or prone to regress than the second step within Schurz's approach. It is not circular because object induction is not justified by applying object induction, but by applying

[11] Suppose that by applying method-induction in our world we correctly determine many different reliable methods. Suppose there are five methods M1–M5 that sometimes make different predictions, but all have a limiting success rate of 0.8. Are all methods M1–M5 justified via method-induction and if so, which of the methods are we supposed to follow? Schurz ([2019]) argues that the following answer is problematic: Always follow the method that is the best at the time. A better answer is: Follow the best method and only switch to another method, if the other method exceeds the old favourite by a certain threshold (see for relevant details Schurz [2019], sec. 6.2). Proponents of strategy (II) are free to adopt the second answer in the light of the problematic case above.

method induction. Furthermore, it does not lead into a regress because the potential regress is stopped by justifying method induction via (i)*–(iv)* a priori. Thus, the familiar problems of attempts to justify induction a posteriori do not concern strategy (II) any more than they concern Schurz's version of strategy (I). Strategy (II) is therefore able to preserve one of the key advantages of Schurz's approach. Is strategy (II) also able to circumvent the problems that we have identified for Schurz's proposal?

In contrast to Schurz's approach, the method-inductive strategy (II) does not rely on modelling scientific methods at the object level as elements in pool Π of a prediction game. This is beneficial for two reasons: first, predictions on the object level do not have to be modelled as values in the range of [0,1].

Second, in contrast to Schurz's account, strategy (II) is also applicable to inductive inferences of general, law-like statements. Schurz's meta-induction cannot be applied to these inferences because inductive methods that aim at law-like generalizations cannot be modelled as elements in pool Π of a prediction game (Arnold [2010], p. 591; Sterkenburg [forthcoming], p. 6). Strategy (II) does not rely on this kind of modelling, so there is no reason strategy (II) should not be applicable to inductive inferences of general statements. In the context of strategy (II), the methods with the best track records on the object level are determined by applying the straight rule. And this is also possible with respect to methods that infer general statements. The method which, in relation to all general, law-like statements made by it, delivers the fewest false ones is the most reliable. Thus, in contrast to Schurz's strategy (I), strategy (II) is also applicable to inductive generalizations. In fact, strategy (II) can be applied to scientific methods in general; for example, it can also be applied to inferences to the best explanation.

A third advantage of the rehabilitation of Reichenbach's vindication offered by strategy (II) is that, in contrast to Schurz's meta-inductive strategy (I), it does not make the justification of object induction seem unmotivated. Just as with meta-induction, method induction is justified to be optimal a priori. However, in contrast to meta-induction, method induction does not make predictions of events at the object level, but only predictions concerning the reliability of prediction methods. Thus, strategy (II) does not make the justification of object induction seem unmotivated.

Fourth, in contrast to the meta-inductive strategy (I), strategy (II) does clearly justify the reliability of the method of object induction and not only its current predictions. It is interesting to note, however, that (II) does not justify that object induction will for all eternity be a reliable method. It is not a priori excluded that by applying method induction at some point, other methods on the object level are judged to be more reliable than object induction. Based on the optimality argument (i)*–(iv)*, scientists are justified in always holding on to method induction, but dogmatically clinging to object induction cannot be justified via strategy (II).

Fifth, in contrast to Schurz's meta-inductive strategy (I), the method-inductive strategy (II) does not depend on complex mathematical considerations in a particular field of computational learning theory (prediction with expert advice). It is therefore less complex and easier to grasp. In times of emerging scepticism toward science in parts of our society, this is an important advantage of strategy (II). Thanks to its relative simplicity, strategy (II) has a better chance of convincing sceptics with respect to science and advocates of nonscientific prediction methods of the superiority of scientific methods over nonscientific ones.[12]

However, it is important to note that besides these advantages, strategy (II) also has some disadvantages compared to Schurz's strategy (I). The first step of strategy (II) is based on the argument (i)*–(iv)*, which is structurally analogous to Reichenbach's original argument (i)–(iv). The new argument (i)*–(v)* avoids the most devastating objection with respect to (i)–(iv)—in contrast to premise (iii), premise (iii)* is not false. However, some of the other problems carry over to (i)*–(iv)*. This is especially worrisome with respect to the problem of being too dependent on the long run and the problem of asymptotic methods (see the second and third problem addressed in section 2).

Schurz's approach does not depend on the long run; within his account, short-run optimality can also be proven for awMI (Schurz [2019], sec. 6.6). This is a decisive advantage of his account. It allows Schurz to deal with methods that show certain feedback effects between predictions and events, and will, therefore, not converge to a stable limit (Schurz [2019], secs. 6.1–6.3).

In addition, Schurz's meta-induction avoids the problem of asymptotic methods. There are many methods that converge the straight rule asymptotically and are, therefore, also optimal in the long run. This generates a serious problem for the second step of strategy (II) because it is unclear whether applying those asymptotic methods now will select object induction at the object level. Strategy (I) avoids this difficulty. Strategy (I) does not depend on the straight rule but on awMI (see section 3). Even though there might be many different meta-inductive methods that converge awMI and are, therefore, also optimal in the long run, only awMI is optimal in the short run. And it is plausible that applying awMI now would in fact favour object induction on the object level. Thus, short-run optimality helps Schurz solve the problem of asymptotic methods and thereby sufficiently restrict what we consider optimal methods, so that they can do what they are supposed to do in the second step of his argument.

The fact that both problems, the problem of being too dependent on the long run and the problem of asymptotic methods, do not arise with respect to strategy (I) is an important advantage of Schurz's account. In contrast, strategy (II) is only successful if these problems can be solved in one of the ways indicated in section 2 (or any other way for that matter).[13]

## 6 Conclusion

With regard to the problem of induction, Reichenbach accepts that induction cannot be justified to be reliable. What he thinks can be justified, however, is that induction is an optimal prediction method: it is guaranteed not to be worse and possibly better than any alternative. Reichenbach's vindication of induction faces various difficulties, one of which is particularly severe (namely the reliable-soothsayer-objection, see the fifth problem in section 2). In the face of this difficulty, two strategies to rehabilitate Reichenbach's account

[12] Note that strategy (II) also has a chance to circumvent the issue raised by Sterkenburg ([2019]) for Schurz's meta-inductive strategy of not being able to deal adequately with an expanding pool of object-level prediction methods, s. fn. 9.

[13] For a thorough discussion and an alternative solution to the problem of asymptotic methods in the context of Reichenbach's original suggestion, see (Salmon [1991], pp. 103–07 and 113–19).

can be differentiated: (I) the meta-inductive strategy (s. section 4), and (II) the method-inductive strategy (s. section 5). Strategy (I) has been motivated, developed, and defended by Schurz ([2008], [2019]). It is undoubtedly a highly interesting and powerful attempt at a justification of induction on the basis of Reichenbach's ideas. In contrast, strategy (II) has not been suggested thus far. I have argued that strategy (II) is also able to rehabilitate Reichenbach's account against its most serious objection and that it is in various respects even superior to Schurz's proposal.

The most important advantages are first that strategy (II) does not rely on mathematical considerations in a particular branch of computational learning theory (prediction with expert advice) and the corresponding need to mathematically model inductive methods as elements of a pool Π in a prediction game. This is an advantage because some of the modelling decisions in this regard are questionable. Second, and related to this, strategy (II) is not restricted to inductive inferences concerning particular events but also applies to inductive inferences concerning general and law-like statements. In fact, it can be applied to all scientific methods, no matter how we want to specify them in detail. Third, strategy (II) is much simpler and easier to grasp. It is therefore a more effective tool to defend science against attacks from pseudoscience and proponents of nonscientific prediction methods.

An additional point that is interesting from an exegetical perspective is that strategy (II) is closer to Reichenbach's own position. Regarding the main difficulty of his original proposal, Reichenbach himself already points out that the inductive straight rule is not only applicable on the object but also on the method level. He does not spell out how exactly this insight is supposed to save his original proposal. But he seems to assume that nothing more than the straight rule and the different levels of its application is needed for this purpose. Strategy (II) illustrates that this assumption is correct.

## Acknowledgements

Thanks to Claus Beisbart, Heiko Briesen, Christian J. Feldbacher-Escamilla, Stefan Fischer, Wolfgang Spohn, Tom F. Sterkenburg, and Verena Wagner for helpful comments on previous drafts of the paper. I would also like to thank Gerhard Schurz for interesting discussions on meta-induction and two anonymous referees for very detailed, helpful and constructive criticism, which led to significant improvements of the paper. This research was supported by the German Research Foundation (Deutsche Forschungsgemeinschaft, grant no. BR 6343/1-1).

*Department of Philosophy*
*University of Konstanz*
*Konstanz, Germany*
*jochen.briesen@uni-konstanz.de*

## References


Arnold, E. [2010]: 'Can the best alternative-justification solve Hume's problem?', *Philosophy of Science*, **77**, pp. 584–93.

BonJour, L. [1998]: *In Defense of Pure Reason: A Rationalist Account of A Priori Justification*, Cambridge: Cambridge University Press

Cesa-Bianchi, N. and Lugosi, G. [2006]: *Prediction, Learning, and Games*, Cambridge: Cambridge University Press.

Van Cleve, James [1984]: 'Reliability, Justification, and Induction', *Midwest Studies in Philosophy*, **9**, pp. 555–67.

Feigl, Herbert [1950]: 'De Principiis non disputandum…? On the meanings and Limits of Justification', in M. Black (ed.), *Philosophical Analysis*, Ithaca, NY: Cornell University Press, pp. 119–56.

Freitag, W. [2015]: 'I Bet You'll Solve Goodman's Riddle', *The Philosophical Quarterly*, **65**, pp. 254–67.

Goodman, N. [1955]: *Fact, Fiction and Forecast*, Cambridge, MA: Harvard University Press.

Henderson, L. [2020]: 'The Problem of Induction', in E. N. Zalta (*ed.*) *The Stanford Encyclopedia of Philosophy*, available at <https://plato.stanford.edu/archives/spr2020/entries/induction-problem/>.

Herz, P. [1936]: 'Kritische Bemerkungen zu Reichenbachs Behandlung des Humeschen Problems', *Erkenntnis*, **6**, pp. 25–31.

Hume, D. [2009]: *A Treatise of Human Nature*, (first published in 1739), Oxford: Oxford University Press.

Hume, D. [2010]: *An Enquiry Concerning Human Understanding*, (first published in 1748), Oxford: Oxford University Press.

Kant, I. [1998]: *Critique of Pure Reason*, (first published in 1781), translated by P. Guyer and A. W. Wood, Cambridge: Cambridge University Press.

Lange, M. [2011]: 'Hume and the Problem of induction', in D. Gabbay, S. Hartmann and J. Woods (*eds.*), *Inductive Logic*, *Handbook of the History of Logic*, *Vol. 10*, Amsterdam: Elsevier, pp. 43–92.

Papineau, D. [1992]: 'Reliabilism, Induction, and Scepticism.' *The Philosophical Quarterly*, **42**, pp. 1–20.

Popper, Karl [2002]: *The Logic of Scientific Discovery*, (first published 1959), London: Routledge.

Reichenbach, H. [2006]. *Experience and Prediction: An Analysis of the Foundations and the Structure of Knowledge*, (first published 1938), Chicago: University of Chicago Press.

Reichenbach, H. [1949]: *The Theory of Probability*, Berkeley/Los Angeles: University of California Press.

Salmon, W. C. [1957]: 'Should we Attempt to Justify Induction?', *Philosophical Studies*, **8**, pp. 45–47.

Salmon, W. C. [1966]: *The Foundations of Scientific Inference*, Pittsburgh: University of Pittsburgh Press.

Salmon, W. C. [1991]: 'Hans Reichenbach's Vindication of Induction', *Erkenntnis*, **35**, pp. 99–122.

Schurz, G. [2008]: 'The Meta-inductivist's Winning Strategy in the Prediction Game: A New Approach to Hume's Problem', *Philosophy of Science*, **75**, pp. 278–305.

Schurz, G. [2018]: 'Optimality Justifications: New Foundations for Foundation-Oriented Epistemology', *Synthese*, **195**, pp. 3877–97.

Schurz, G. [2019]: *Hume's Problem Solved: The Optimality of Meta-Induction*, Cambridge, MA.: Massachusetts Institute of Technology Press.

Schurz, G. [forthcoming]: 'Optimality Justifications and the Optimality Principle: New Tools for Foundation-Theoretic Epistemology', *Nous*, available at <https://doi.org/10.1111/nous.12390>.

Skyrms, B. [1965]: 'On Failing to Vindicate Induction', *Philosophy of Science,* **32**, pp. 253–68.

Skyrms, B. [1975]: *Choice and Chance: An Introduction to Inductive Logic*, Encino, CA: Dickenson.

Sterkenburg, T. F. [2019]: 'The Metainductive Justification of Induction: The Pool of Strategies', *Philosophy of Science,* **86**, 981–92.

Sterkenburg, T. F. [2020]: 'The Meta-Inductive Justification of Induction', *Episteme*, **17**, pp. 519–41.

Sterkenburg, T. F. [forthcoming]: 'On Explaining the Success of Induction', *The British Journal for the Philosophy of Science*, available at <https://doi.org/10.1086/717068 >.

Zinke, A. [2020]: 'Against Grue Mysteries', *Erkenntnis*, **85**, pp. 1023–33.