DOI: 10.1017/psa.2025.10130
This is a manuscript accepted for publication in *Philosophy of Science*.
This version may be subject to change during the production process.

Identification and Indetermination in the Meta-Inductive Approach to Induction

Kabir S. Bakshi [1]

[1] History and Philosophy of Science, University of Pittsburgh
Corresponding author: Kabir S. Bakshi; kabir.bakshi@pitt.edu

Abstract
The meta-inductive approach to induction justifies induction by proving its optimality. The argument for the optimality of induction proceeds in two steps. The first "a priori" step intends to show that meta-induction is optimal and the second "a posteriori" step intends to show that meta-induction selects object-induction in our world. I critically evaluate the second-step and raise two problems: the identification problem and the indetermination problem. In light of these problems, I assess the prospects of any meta-inductive approach to induction.

1. Introduction

My aim here is to critically engage with the novel meta-inductive approach to induction. Broadly speaking, the meta-inductive approach provides a justification of induction through a two step process. The first "a priori" step consists in proving that meta-induction is optimal and the second "a posteriori" step consists in showing that meta-induction selects object-induction in our world.

Current engagement with and discussion of the meta-inductive approach have mostly been focused on the first-step.[1] Departing from the existing literature, I will here focus on the second-step. Although the first-step faces some serious challenges (expanding pool of strategies, non-convex loss functions, appeal to bounded cognition, etc.), for the purposes of this paper I grant the the first-step. Even with this substantial concession, I show that the meta-inductive approach faces significant obstacles.

I proceed as follows. In section 2, I briefly provide a sketch of the meta-inductive approach. This will help make clear the bipartite structure of the meta-inductive justification. Following this, I develop the *identification problem* (section 3) and the *indetermination problem* (section 4) for the meta-inductive approach. Section 5 concludes.

[1] For example, Arnold (2010) provides reasons to be skeptical of the notion of optimality in an infinite setting and Sterkenburg (2019) critiques the meta-inductive approach for failing to deal with an expanding pool in the first-step. See also Douven (2023) for some comments on the meta-inductive approach.

	Uniform State	Non-Uniform State
Induction	Successful	Non-Successful
Non-induction	Successful or non-successful	Non-Successful

Figure 1. Reichenbach's matrix. As reconstructed in Salmon (1974).

A preliminary before I start. I use 'meta-induction' and 'object-induction' to distinguish between induction at the level of methods (meta-inductions) and induction at the level of objects (object-induction). However, where the context admits I sometimes use 'induction' to stand for object-induction (but never for meta-induction).

2. The meta-inductive approach to induction

The pragmatic approach to induction – defended most famously by Hans Reichenbach – accepts that in light of Hume's skeptical challenge a justification for the reliability of induction is not possible (Reichenbach, 1938, 1949). Instead, Reichenbach suggests that a more modest aim is achievable: we can justify induction not by showing its reliability but by showing its optimality. On one reading of Reichenbach's proposal, induction is justified because it is the most optimal method in predicting finite frequencies of observations.[2] He argues for this claim using an argument by cases. Dividing the space of states of the world into uniform and non-uniform states, he argues as follows (cf. Figure 1):

- If the state of the world is uniform then the world is induction friendly. Hence, induction will be the best method vis-à-vis predictive success. Other methods might be successful but induction will be at least as successful as any other method. Thus induction is optimal (compared to all the other methods) over uniform states of the world.
- If, however, the state of the world is non-uniform then induction will not be predictively successful. But in such a case no other prediction method will be successful as well. Thus, induction is optimal over non-uniform states of the world.

The point of weakness in Reichenbach's argument is his claim that if the state of the world is non-uniform then no non-inductive method can be successful. As Salmon (1974) points out, this is incorrect. Consider a non-uniform state of the world where an oracle has foreknowledge. Even though the world is inhospitable to induction, it is a world where a non-inductive method can have an arbitrarily high success rate. Hence, Reichenbach's claim that induction is optimal when the state of the world is non-uniform fails. Facing this objection, Reichenbach tweaked his account to explicitly take into consideration situations where a non-inductive method can outperform induction. He argued that in such cases (for *e.g.* oracle in a non-uniform world) induction will still be as successful as the oracle because an inductive method will latch

[2] Schurz (2021) distinguishes three versions of Reichenbach's argument and traces the historical trajectory of their development.

https://doi.org/10.1017/psa.2025.10130 Published online by Cambridge University Press

onto the oracle's predictions. Reichenbach, however, provides no argument to support the claim that such a strategy can be implemented. Taking the advice from Salmon that Reichenbach provides "a valid core from which we may attempt to develop a more satisfactory justification" of induction, the meta-inductive approach to induction aims to provide a rigorous argument to show just how induction is optimal (Salmon, 1967, 54).

The meta-inductive approach uses tools and results from machine learning, in particular from the area of online learning under expert advice (Cesa-Bianchi and Lugosi, 2006; Rieskamp and Otto, 2006; Dieckmann and Todd, 2012). The framework used is that of prediction games and consists of the following components:

(1) *Events*
Individual events are the target for predictions and a stream of events is represented by an infinite sequence of events $\mathbf{e} = e_1, e_2, e_3, \ldots$. Each individual event is represented by a real number in the interval $[0, 1]$ and can be thought of as a measurement or observation of the quantity of interest.[3] Denote by e_n the n-th event in an event sequence $\mathbf{e}$.

(2) *Methods*
Every non-meta-inductive method (or player or prediction strategy) p is a member of the *finite* pool of method $\Pi = \{p_1, p_2, \ldots, p_n\}$.[4] Each prediction strategy specifies a prediction for the next event. Denote by $\Pi{*} = \{p_1, p_2, \ldots, p_n, \mathsf{MI}\}$ the pool of methods in Π and the meta-inductive method MI.

(3) *Prediction*
A prediction $\mathscr{P}_n^p$ at stage n for the event at stage $n+1$ by a method $p \in \Pi{*}$ is a real number in the interval $[0, 1]$.[5]

(4) *Loss Function*
A loss function $\mathscr{L}(\mathscr{P}_n^p, e_n)$ is a measure of how far the prediction of method p was from the true value of event e_n at stage n of the prediction game. Loss functions are monotonic and convex.[6]

(5) *Prediction Games*
A prediction game $(\mathbf{e}, \Pi{*}, \mathscr{L})$ is a tuple of a sequence of events, a pool of methods, and a loss function. At each stage of the game, methods in the pool make predictions for the next stage of the game. These predictions are then evaluated with respect to the true value of the event using the loss function. One can then define the notion of

[3] If the range of the measurement of the quantity is $-b$ to a, then one can always normalize by a measurement c by first subtracting $-b$ and then multiplying by $1/a+b$.

[4] The "realistic assumption" (Schurz, 2019, 50) of restricting prediction games to finitely many predictors has come under scrutiny Arnold (2010); Sterkenburg (2019).

[5] There are some qualifications here with discrete predictions and real-valued predictions that are not important for present purposes (Schurz, 2019, 147 ff.).

[6] A loss function is convex just in case the loss of a weighted mean of two predictions is not more than the sum of the weighted mean of the losses of those two predictions. Although Schurz proves his results using concave loss functions, he also briefly considers convex loss functions (Schurz, 2019, 156).

https://doi.org/10.1017/psa.2025.10130 Published online by Cambridge University Press

the *success rate* of each method in the pool. The success rate of a method p at the n-th stage of the prediction game is

$$\mathsf{succ}_n(p) = \frac{1}{n}\sum_{i=1}^{n-1}\left(1 - \mathscr{L}(\mathscr{P}_i^p, e_i)\right). \tag{1}$$

A meta-inductive method, MI, is a method which predicts at stage n by accessing the prediction of all the other non-inductive methods at n. The simplest meta-inductive method – "imitate the best" – predicts at n identical to the non-inductive method with the highest success rate at n.[7] Schurz (2019) discusses other more complex meta-inductive methods including the attractivity-weighted method, the exponentially-weighted method, and the collective method. I will use the result of the attractivity-weighted method ("aMI") below.[8]

The meta-inductive justification of induction consists of two steps. The goal of the first step is to show that the meta-induction – induction over prediction methods – is *universally access optimal*. A method p is *universally access optimal* just in case it is optimal in all prediction games where the methods in $\Pi - \{p\}$ are accessible to p. A method $p' \in \Pi - \{p\}$ is *accessible* to p iff p knows the prediction made by p' at any stage of the prediction game and keeps a record of p' success rate in memory. The notion of universal access optimality is operationalized in the meta-inductive framework by an attractivity (or equivalently a regret) parameter which measures how much a method is attractive to the MI method. For a MI, the attractivity with respect to a method p at a stage n is:

$$\mathsf{Att}_n(p) = \mathsf{succ}_n(p) - \mathsf{succ}_n(\mathsf{MI}) \tag{2}$$

With this in hand, the attractivity-weighted meta-inductive (aMI) method can be formulated through the following prediction strategy:

$$\mathscr{P}_{n+1}^{\mathsf{aMI}} = \frac{\sum_{p\in\Pi} w_n(p)\cdot\mathscr{P}_{n+1}^p}{\sum_{p\in\Pi} w_n(p)} \tag{3}$$

where $w_n(p)$ is the factor by which aMI weighs the prediction of p at stage n using the attractivity:

$$w_n(p) = \begin{cases} 0 & \text{if } \mathsf{Att}_n(p) < 0 \\ \mathsf{Att}_n(p) & \text{otherwise} \end{cases} \tag{4}$$

It can then be proved that for all-convex loss functions and for all $p \in \Pi$ and for all $n \geq 1$ (Schurz, 2019, 143):

$$\mathsf{succ}_n(\mathsf{aMI}) - max(\mathsf{succ}_n(p)) \geq \sqrt{\frac{m}{n}} \tag{5}$$

[7] The imitate the best method has also been studied in Schurz (2008) "simple meta-inductivist"; Cesa-Bianchi and Lugosi (2006) "follow-the-best-expert"; and De Rooij et al. (2014) "follow-the-leader".

[8] The aMI has also been studied as "the weighted-majority algorithm" in Littlestone and Warmuth (1994); Cesa-Bianchi and Lugosi (2006); and Shalev-Shwartz and Ben-David (2014).

https://doi.org/10.1017/psa.2025.10130 Published online by Cambridge University Press

where m is a constant which depends on the size of Π. In the long-run limit the loss rate of aMI is strictly no less worse than any other method. That is

$$lim_{n\to\infty}(\mathsf{succ}_n(\mathsf{aMI}) - max(\mathsf{succ}_n(p))) = 0 \quad (6)$$

This converge result (and the short-run result) is taken to show that meta-induction is universally access optimal. The proponent of the meta-inductive approach thus concludes that if prediction is our goal, then in a very rigorous and clear sense, we are best-off with meta-induction (Schurz, 2019, 199). Sure, the proponent concedes, some states of the world will be such where an oracle or a soothsayer will predict perfectly. But without making any assumptions about the state of the world, meta-induction will be the best strategy to follow.

But we should not miss the woods for the trees. The project is to justify object-induction not meta-induction. The first-step in the argument is an argument for the universal access optimality of meta-induction and the second-step is supposed to provide a justification of object-induction and an answer to Hume's challenge. The second-step establishes the claim that the predictions of meta-induction are identical (or arbitrary close) to object-induction in our actual world. It is interesting to note, however, that Schurz dedicates less than two pages in his 372 page tome in discussing the second-step of the argument. I quote his argument for the second-step in full:

> Until the present time and according to the presently available evidence, object-inductive methods dominated noninductive methods in the following sense: in many fields some object-inductive method was significantly more successful than every noninductive method, though in no field was a noninductive method significantly more successful than all object-inductive methods. (Schurz, 2019, 209)

The observation that no non-inductive method has been substantially more successful than any object-inductive method forms the basis of the meta-inductive conclusion that the universal access optimality of meta-induction provides an a posteriori justification of object-induction (Schurz, 2021, 987). Because we have evidence for the significant success of object-induction, the meta-inductive argument goes, meta-induction selects object-induction in our world. And since we are justified in following meta-induction because of the first-step of the argument, we are thus justified in following object-induction. The warrant, in this sense, flows from the level of methods where meta-induction lives to the level of object-induction. Thus, on the meta-inductive approach, this two step strategy provides a non-circular justification of object-induction, solving Hume's challenge (Schurz, 2019, 197).

3. Identification in Meta-Induction

Meta-induction selects object-induction, and hence justifies it, only if object-induction has so far been the most successful method. But how is object-induction identified? This is the *identification problem* for the meta-inductive approach to induction.

https://doi.org/10.1017/psa.2025.10130 Published online by Cambridge University Press

A first-pass response is to claim that it is futile to identify object-induction generally. Instead we must identify object-induction by looking in the domains where it has been successful and unsuccessful. The motivation for this is the observation that object-induction has been extremely successful in many domains of scientific and non-scientific inquiry while simultaneously being unsuccessful in other domains. On the one hand, for example, object-induction predicts that the sun will rise tomorrow. Similarly object-induction has been significantly successful in scientific inquiry. For example, by induction on the mass of observed electrons, object-induction informs us that all electrons will have a mass of $9.11 * 10^{-31}$ kg. On the other hand, object-induction has been unsuccessful in many domains of scientific and non-scientific inquiry as well. For example, object-induction provides little useful guidance in predicting stock prices or in guiding us in the area of physics beyond the standard model.

The first-pass response solves the identification problem by a divide-and-conquer approach: we should only investigate induction-friendly domains. The method which works in those domains is induction. However, this divide-and-conquer response to the identification problem faces difficulties. Sure, there are induction hostile and induction hospitable domains but this response does not give any guidance in picking out object-induction from among other methods in the pool of methods. The identification problem bestows a duty on the meta-inductive approach to provide a principled way to identify object-induction. Only once object-induction is identified can its track record be checked and (if it is significantly successful than other methods) object-induction justified. Until object-induction can be identified, Hume's challenge remains.

A better response than the first-pass divide-and-conquer response was recently proposed in Schurz and Thorn (2020). According to them, Norton's material theory of induction provides an elegant criterion to identify object-inductive methods in a specific domain (Norton, 2003, 2021). On the material theory of induction, there is no universal schema of induction. Instead, all induction is local and tied to a domain. Norton's project, as he emphasises, is not to provide an answer to Hume's challenge but to provide a criterion to distinguish between good and bad scientific inductive inferences. In the material theory, an induction is good just in case there are local background facts which warrant the particular induction. Norton uses the analogy of placing stones in a self-supporting arch to illuminate the structure the local background facts must conform to. Each stone or background fact is independently well confirmed and supports every other fact in the arch.

Schurz and Thorn argue that the material theory and the meta-inductive approach are "complementary" (Schurz and Thorn, 2020, 93). The material theory solves the identification problem for meta-induction by supplying "domain-specific aspects of object-induction". The thought is that the material theory provides a principled criterion to identify induction-friendly domains and pick out object-induction. Of course, because the material theory denies that any universal schema of induction exists, the method the material theory specifies will be local and domain dependent. But it is claimed that that is no bug but a feature.

https://doi.org/10.1017/psa.2025.10130 Published online by Cambridge University Press

> ...there are many domains in which the superiority of object-inductive prediction methods is not obvious. This latter point brings us back to Norton's material account of object-induction. It is only in domains that are regulated by strong uniformities that the superiority of object-inductive methods over non-inductive methods of prediction will be strong enough that it can convince even skeptical persons. The strength of Norton's material account of object-induction lies in the fact of illuminating the detailed structure of these local uniformities that make the success of inductions in science possible. (Schurz and Thorn, 2020, 92)

This might seem to allay the problem for meta-induction raised by the identification problem. However, this is mistaken. The material theory and meta-induction are in tension. The material theory and meta-induction are not complementary. Here are two reasons why.

First. The clearest way to see the tension inherent between the material and the meta-inductive approach is by noting that there are cases where their judgments diverge.

Examples of this kind are abundant because of the disparate weights meta-induction and material theory place on past successes. While which (if any) method is picked out on the meta-inductive approach depends entirely on the track record of the method, the material theory has no time for past successes. The material theory, unlike the meta-inductive approach, is 'memory-less': the success or failure of past predictions do not play any role in judging whether an inductive inference is good or bad. In the material theory, it is the background facts of the domain and not facts about the success of the methods which determine the goodness of an inductive inference.

One can manufacture cases where the background facts obtain for an inductive inference but the success rate of any object-inference method has not yet been substantially greater than other methods in the pool. In such cases the material theory will judge inductive inferences as good while the meta-inductive approach will not deem object-induction as the favoured method.

Consider as an example of such a case the discovery of the properties of radium chloride by Marie Curie, discussed in Norton (2021). According to the material theory, Curie's inference that radium chloride has the same crystalline form as barium chloride is a good inductive inference. It is good because it is warranted by a local background fact – the *Weakened Haüy's Principle*. But meta-induction does not pick out Curie's inference because at the turn of the century it was no where near (substantially) more successful than any other predictive method.

Indeed, many scientific discoveries only become significantly successful after some time has elapsed from their conception. Because it does not depend on the success rate of a method, the material theory will still judge these inferences to be good. However, the meta-inductive approach will not judge these inferences as the ones which we should follow. Such examples can be easily multiplied.

https://doi.org/10.1017/psa.2025.10130 Published online by Cambridge University Press

Second. The other way in which the material theory and the meta-inductive approach are in tension relates to a crucial background assumption. A central assumption of the meta-inductive approach is the assumption of cognitive finiteness (Schurz, 2019, 266). Schurz uses this assumption to defend the meta-inductive approach against Arnold (2010)'s critique of the finite pool of methods. According to the cognitive finiteness assumption, the fact that humans are cognitively bounded agents play an important and central role in the justification of induction. For example, the finiteness of Π is justified because the cognitive finiteness assumption precludes having infinite methods under consideration.

Using such an assumption to justify induction is an anathema to the material-theorist. In the material account of induction, there is no place *at all* for considerations about human cognition and capabilities. Indeed, the material account goes even further. Humans or any other inference doers are not necessary for the material account. For example, if humans did not exists, there would still be good and bad inductive inference. From the viewpoint of the meta-inductive approach which justifies object-induction based on the success of object-induction, such a claim seems bizzare. On the meta-inductive approach, if there are no inference doers there is no method one should follow.

Thus the appeal to the material theory to assuage the identification problem fails because of inherent and inescapable tension between the material theory and the meta-inductive approach. The identification problem remains a serious issue for the meta-inductive approach.

4. Indetermination in Meta-Induction

Suppose, modulo the identification problem presented above, there is a principled way to identify object-induction. Even so, the second-step fails because meta-induction fails to *uniquely* select object-induction. This is the *indetermination problem* for the meta-inductive approach.

Indetermination arises because – at any stage – there exist many methods which have the success rate (and hence the same attractivity to a meta-inductive method) as object-induction. And moreover, these equal-success-rate methods diverge from object-induction at some stage in the future and are non-identical to object-induction. That this is the case is straightforward to see. Recall that event sequences consist of infinite individual events $\mathbf{e}$ (where $\mathbf{e} = e_1, e_2, e_3, \ldots$). Now consider the set $\Pi_G \subset \Pi$ of *Goodman methods* at the stage n. The membership of Π_G at n can be enumerated as follows:

p_1: method p_1 predicts identical to object-induction up until n but predicts differently from object-induction at $n+1$.
p_2: method p_2 predicts identical to object-induction up until n but predicts differently from object-induction at $n+2$.
p_3: method p_3 predicts identical to object-induction up until n but predicts differently from object-induction at $n+3$.
.
.

.

Even if one grants the assumption that Π contains finitely many methods, the indetermination problem remains nebulous. The second-step of meta-induction does not provide a justification for object-induction uniquely. At best, if everything works, at n the second-step selects all the methods in Π_G plus object-induction. The meta-inductive approach cannot justify induction *uniquely*.

This result in significantly more problematic for the meta-induction than the challange presented in Sterkenburg (2020). Sterkenburg argues that meta-induction cannot provide justification for object-induction simpliciter. Instead he claims that the second-step of the argument warrants a more modest claim: object-induction is justified *now*. However, indetermination using Goodman methods blocks even this more modest claim. Given the indetermination problem, meta-induction cannot even justify object-induction now. The claim must be even more modest. The meta-inductive approach provides us, at most, that some method in Π_G is justified now. From this very weak conclusion, it does not follow that object-induction is justified. This is very far from responding to Hume's challenge.

Is there a response to the indetermination highlighted here? A potential reply is provided by Schurz in discussing what he calls the *selection problem*, the problem of selecting which methods to include in Π. In providing a criterion for selecting methods that go into the pool Π, Schurz argues against the inclusion of Goodman methods by pointing out that they are algorithmically more complex than non-Goodman methods. He appeals to the notion of Kolmogorov complexity to choose methods which have higher "epistemic naturalness" (Schurz, 2019, 269). Appealing to complexity to exclude Goodman methods from entering into Π from the start trivially dissolves the indetermination problem. However, this reply does not work.

If the meta-inductive approach includes any inductive assumption, charges of circularity will naturally follow and the meta-inductive justification of induction will fail. This is evident in Schurz's rejection of Goodman's solution of inductive projectability to the the new problem of induction on the basis that Goodman's solution is "circular" (Schurz, 2019, 271). But in appealing to algorithmic simplicity and epistemic naturalness to solve the selection problem (and the indetermination problem), the meta-inductive inadvertently makes an inductive (and hence an inadmissible) assumption.

To claim that one should favor simpler rather than complex methods results is making a claim about the world. Thinking that only simpler non-Goodman methods can be included in the pool of methods amounts to weighing more complex Goodman methods ($p \in \Pi_G$) by zero. But that is just to say that we assume that a Goodman method cannot predict – with a perfect success rate – the true event sequence. This is clearly an inductive assumption. If one of the most attractive features of the meta-inductive approach is its independence from any inductive assumptions, then the problem of indetermination *cannot* be solved by any appeal to simplicity (or naturalness).

https://doi.org/10.1017/psa.2025.10130 Published online by Cambridge University Press

Although it true that we observe in our actual worlds that simple non-Goodman methods are more successful than gerrymandered Goodman methods, such an observation cannot support the meta-inductive approach. It cannot solve the indetermination problem. The prospects for the meta-inductive approach look bleak.

5. Conclusion

The meta-inductive approach is a novel and a well thought-out response to Hume's challenge. Two developments are particularly welcomed. First, shifting the focus from the reliability of induction to its optimality seems to be a more fruitful and fertile philosophical program. And second, justifying induction in a two-step process with an empirical second-step is promising because it justifies induction only for our actual world rather than justifying induction for all worlds (something which has proven to be extremely difficult).

Although, the meta-inductive approach should be commended for these aspects, I have raised two problems worries for the approach here. Instead of focusing on the first-step of the argument which has received the most critical attention, I focused on the second step. I introduced two problems for meta-induction: the identification problem and the indetermination problem. Both are serious and possibly fatal to the meta-inductive approach. I entertained some potential replies but concluded that none are helpful. These problems are especially concerning for any meta-inductive approach because they block the empirical part of the approach and thus these cannot be resolved by theoretical advances in machine learning which is the foundation for the meta-inductive approach.

Acknowledgments. For her discussions, encouragement, and comments on previous versions of this paper, I am grateful to Francesca Zaffora Blando. Thanks also to John D. Norton and audiences at the 29th PSA biennial in New Orleans and the Pittsburgh History and Philosophy of Science WiP seminar in Fall 2024.

Conflicts of Interest. None to declare.

Funding Disclosure. None to declare.

References

Arnold, E. (2010) Can the best-alternative justification solve Hume's problem? On the limits of a promising approach. *Philosophy of Science*. *77*(4), 584–593. https://doi.org/10.1086/656010.

Cesa-Bianchi, N. and Lugosi, G. (2006) *Prediction, Learning, and Games*. Cambridge University Press. https://doi.org/10.1017/CBO9780511546921.

De Rooij, S., Van Erven, T., Grünwald, P. D. and Koolen, W. M. (2014) Follow the leader if you can, hedge if you must. *The Journal of Machine Learning Research*. *15*(1), 1281–1316. https://dl.acm.org/doi/10.5555/2627435.2638576.

Dieckmann, A. and Todd, P. M. (2012) Simple rules for ordering cues in one-reason decision making In *Ecological rationality: Intelligence in the world*, Todd, P. M. and Gigerenzer, G. (eds). Oxford University Press. pp. 274–306. https://doi.org/10.1093/acprof:oso/9780195315448.003.0078.

Douven, I. (2023) Explaining the success of induction. *British Journal for the Philosophy of Science*. *74*(2), 381–404. https://doi.org/10.1086/714796.

Littlestone, N. and Warmuth, M. (1994) The weighted majority algorithm. *Information and Computation*. *108*(2), 212–261. https://doi.org/10.1006/inco.1994.1009.

Norton, J. D. (2003) A material theory of induction. *Philosophy of Science*. *70*(4), 647–670. https://doi.org/10.1086/378858.

Norton, J. D. (2021) *The Material Theory of Induction*. University of Calgary Press. Calgary, Alberta, Canada. https://doi.org/10.2307/j.ctv25wxcb5.

Reichenbach, H. (1938) *Experience and Prediction: An Analysis of the Foundations and the Structure of Knowledge*. University of Chicago Press. Chicago, IL, USA. https://doi.org/10.1037/11656-000.

Reichenbach, H. (1949) *The Theory of Probability*. University of California Press. Berkeley, CA.

Rieskamp, J. and Otto, P. E. (2006) SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General*. *135*(2), 207–236. https://doi.org/10.1037/0096-3445.135.2.207.

Salmon, W. C. (1967) *The Foundations of Scientific Inference*. University of Pittsburgh Press. Pittsburgh, PA.

Salmon, W. C. (1974) The pragmatic justification of induction In *The Justification of Induction*, Swinburne, R. (eds). Oxford University Press. pp. 85–97.

Schurz, G. (2008) The meta-inductivist's winning strategy in the prediction game: A new approach to hume?s problem. *Philosophy of Science*. *75*(3), 278–305. https://doi.org/10.1086/592550.

Schurz, G. (2019) *Hume's Problem Solved: The Optimality of Meta-Induction*. The MIT Press. Cambridge, Massachusetts. https://doi.org/10.7551/mitpress/11964.001.0001.

Schurz, G. (2021) Reichenbach's best alternative account to the problem of induction. *Synthese*. *199*(3-4), 10827–10838. https://doi.org/10.1007/s11229-021-03269-3.

Schurz, G. and Thorn, P. (2020) The material theory of object-induction and the universal optimality of meta-induction: Two complementary accounts. *Studies in History and Philosophy of Science Part A*. *82*, 88–93. https://doi.org/10.1016/j.shpsa.2019.11.001.

Shalev-Shwartz, S. and Ben-David, S. (2014) *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press. https://doi.org/10.1017/CBO9781107298019.

Sterkenburg, T. F. (2019) The metainductive justification of induction: The pool of strategies. *Philosophy of Science*. *86*(5), 981–992. https://doi.org/10.1086/705526.

Sterkenburg, T. F. (2020) The meta-inductive justification of induction. *Episteme*. *17*(4), 519–541. https://doi.org/10.1017/epi.2018.52.