

Review of Gerhard Schurz *Optimality Justifications*, 2024, Oxford: OUP

Richard Pettigrew
Richard.Pettigrew@bristol.ac.uk

April 9, 2024

Optimality Justifications is the triumphant culmination of a research programme pursued by Gerhard Schurz for a little over fifteen years. At its heart is a mathematical result that Schurz proves, building on a tradition of related results from the computational theory of learning.¹ I'll describe this result below. Upon this result, Schurz wishes to build a novel a posteriori justification of inductive inference as a rational method by which to form empirical beliefs. And around this justification, he wishes to construct a comprehensive internalist foundationalist epistemology. The basic beliefs of this foundationalist system are certain analytic truths and certain of those beliefs formed by introspection; and the inferences by which we are justified in forming new beliefs on the basis of ones already justified are classical logical deduction, induction, and abduction (or inference to the best explanation).

Chapters 1 and 2 set out the book's ambition and explain Schurz's reasons for seeking the sort of epistemology he presents. Chapter 3 argues that analytic and introspective beliefs are indeed basic, and Chapter 4 explains how justification transfers across an inference. Chapter 5 describes the novel style of argument that gives the book its title, that is, optimality justifications, which aim to justify using certain sorts of inference. The remainder of the book provides optimality justifications for induction (Chapters 6 and 7), classical logical deduction (Chapter 8), and abduction (Chapters 9 and 10). I'll describe the central instance of this style of argument below—the justification of inductive inference presented in Chapter 6—but I'd first like to note two further things about the book's outlook and ambitions.

The book takes the Quinean view that epistemology is akin to engineering (Quine, 1998). That is, it takes a teleological view of the subject: we fix the goal of beliefs, and the epistemologist's job is to identify those ways

¹Schurz suggests (Cesa-Bianchi & Lugosi, 2006) for an overview of the topic to which his new results contribute.

of forming beliefs that are both feasible for creatures like us and optimal among those that are feasible as ways of achieving the goal in question. For Schurz, a veritist, the goal of belief is truth, an external criterion; but, as mentioned above, Schurz wishes to produce an internalist epistemology, and so, in order to justify basic beliefs or methods of inference that preserve justification, it is not enough to observe that the basic beliefs are true nor that the rules of inference are conditionally reliable, as an externalist veritist like Alvin Goldman might. It must be possible to apprehend from the inside whatever makes them optimal means to the external end of truth; the justifications given must be accessible and persuasive to the individual for whom we're providing them.

In this respect, Schurz's project is analogous to Jim Joyce's (1999) accuracy-based argument for Probabilism, the norm that says your credences should be probabilistic. Probabilism is an internalist norm: whether or not you have probabilistic credences is something you can, in principle, discern from the inside. Joyce's justification appeals to the external criterion of proximity to the truth: he shows that, if you fail to satisfy Probabilism, there are alternative credences you could have had that are guaranteed to be more accurate. But, because the alternatives are guaranteed to be better, this is something you can discern from the inside, without needing any further access to the external world. As we will see below, the analogy is strengthened by the fact that Schurz also appeals to a dominance argument to justify induction.

The internalism is crucial for the overarching ambition of Schurz's project. He is motivated not only to provide a justification for our ordinary and scientific ways of forming beliefs. He seeks an epistemology that allows us to resolve disagreements in order to provide a secure set of empirical beliefs on which we can base our political decisions and discussions—he cites Jürgen Habermas (1984) as an important inspiration. And for this we need an internalist account.

This, then, is the structure, outlook, and ambition of the book. But to appreciate what really makes it novel and worthy of serious philosophical engagement, I must describe what Schurz means by an optimality justification. Though he provides such justifications for classical logical deduction and abduction, I'll describe the optimality justification for inductive reasoning, which is where this project began (Schurz, 2008).

Schurz argues that the predictions of unobserved numerical quantities formed on the basis of enumerative induction from observed cases are justified, at least when forming estimates in such a way has a track record of success in the past. That is, he doesn't argue that enumerative induction is always rational: there are perhaps domains of scientific study in which it has proved unreliable in the past, and applying it in the future in such domains is not justified. But there are domains in which it has proved reliable, Schurz claims, and when dealing with such a domain, continuing to

use it is justified.

Such a style of argument will recall part of Hume's sceptical argument in the *Treatise* for exactly the opposite conclusion: we cannot hope to infer the future success of induction from its past success without using induction itself and thereby begging the question at issue. Schurz agrees, but he argues that, in order to justify using a method of inference to form our judgments, we needn't establish that it is reliable. We need only show that it is in some sense optimally reliable among the various methods of inference available to us—it's the best we can do. After all, rationality requires no more than that: to be rational is to do the best you can within the constraints imposed on you.

At this point, we recall another famous argument concerning induction, namely, Reichenbach's (1938) anti-sceptical argument. Reichenbach claims that, while we cannot be sure induction is reliable, we can know a priori that, if anything is reliable, induction is. This gives a sort of dominance argument for induction: in worlds in which there is a reliable method of inference, induction is reliable; in worlds in which there is no reliable method of inference, nothing is. And so, Reichenbach concludes, we are justified in using induction.

As Schurz notes, it's long been known that Reichenbach's argument fails: there are methods of inference that succeed in worlds in which induction fails. And indeed the so-called 'no free lunch' theorems of Wolpert & Macready (1997) make this observation precise and greatly strengthen it, showing roughly that, in a particular formal representation of the problem, average performance over all possibilities is the same for all inference methods.

Nonetheless, Schurz argues, it's possible to run an a priori argument inspired by Reichenbach's at the meta-level rather than the object-level to show that a meta-method of hewing close to inference methods that have been successful in the past is optimal. And then we can appeal to the past track record of induction to argue that the optimality of this meta-method shows that hewing close to induction is justified. It's an ingenious argument built on Schurz's novel mathematical result that I mentioned above.

To present this result and the argument Schurz builds around it, let me introduce the formal framework in which he considers the problem of induction. We represent the part of the world about which are reasoning by an infinite sequence of real numbers, such as (14, 15, 16, 15, ...). Perhaps this lists the average temperature in degrees Celsius for each season in a particular city; perhaps it lists the number of mountain hares observed during successive months on a particular Scottish hillside. We then represent a method of inference as a function that takes a finite initial segment of such an infinite sequence and returns its prediction for the next number in that sequence. So, for instance, our inference method might take the sequence of average seasonal temperatures observed so far and predict the

next season's to be that.

Now, Schurz supposes, there is a set of inference methods that are cognitively accessible to you—we'll call those the *original methods*. Given the sequence that represents the world and given an inference method, we can generate the sequence of predictions the method makes when successively fed the initial sequences of the world's data, and we can then judge the accuracy of those predictions as a measure of how close the predictions lie to the actual facts.

As well as these original object-level methods, there are also meta-methods. These are functions that take not just an initial segment of the world's data sequence but also the track records of the various object-level methods as given by their accuracy scores and possibly the track records of the meta-method itself and returns a prediction of the next piece of data in the sequence. So a meta-method builds a new object-level method out of the observations so far and the predictions of the various object-level methods.

A natural sort of meta-method is what Schurz calls an *attractivity-weighted method*. The prediction for the next data point given by the new object-level method it generates is a weighted average of the predictions made by the original methods, where the weight given to a method's new prediction is greater the more extra success it has had so far over and above the success that the attractivity-weighted method itself has had so far. It is about this sort of method that Schurz proves his remarkable result. Regardless of the world, regardless of the original object-level methods, and for many natural ways of measuring the accuracy of predictions, the new object-level method created by this meta-method has the following two features: (i) in the long run, the average accuracy of this method, taken over the whole sequence of predictions it makes up to a given point in time, is guaranteed to converge to the maximum average accuracy attained by the original methods up to that time; and (ii) in the short run, there is a bound on how far the average accuracy of this method can lie below the maximum average accuracy attained by the original methods up to that point, and that bound decreases steadily to zero as time progresses.

On the basis of this result, Schurz gives what he considers a dominance argument for using an accuracy-weighted meta-method to construct your object-level method. The goal of a method, Schurz says, is its long-run predictive accuracy. His theorem shows that, relative to that goal, using the object-level method created by the accuracy-weighted meta-method weakly dominates using one of the original methods. In all situations, it's as good, and in some cases it's better. And for this reason, Schurz submits, it's justified. Now, suppose we are working in an area of science in which induction has been the most successful predictor in the past; then, if we wish to predict the future, we are justified in using a weighted average of the originals methods that gives a lot of weight to induction, and so will hew closely to its predictions. Thus, inferring by induction in that area is

justified.

This argument for induction in situations in which it has been successful in the past is the prototype of Schurz's novel method of optimality justifications. The remainder of the book is devoted to providing justifications of this same sort for the other central inference methods that we use.

Optimality Justifications is a book that greatly rewards careful study. It is the product of many years of careful, rigorous, and original thought about one of the deepest problems in philosophy, namely, the normative standing of our everyday and scientific beliefs, and every chapter contains novel insights. Moreover, Schurz's treatment of the theories he ultimately wishes to usurp are charitable, careful, and illuminating. An exemplary piece of philosophy.

References

- Cesa-Bianchi, Nicolò, & Lugosi, Gábor (2006). *Prediction, Learning, and Games*. Cambridge, UK: Cambridge University Press.
- Habermas, Jürgen (1984). *Theory of Communicative Action*, vol. 1. Boston: Beacon Press.
- Joyce, James M. (1999). *The Foundations of Causal Decision Theory*. Cambridge Studies in Probability, Induction, and Decision Theory. Cambridge: Cambridge University Press.
- Quine, W. V. O. (1998). Reply to Morton White. In L. Hahn, & P. A. Schlipp (Eds.) *The Philosophy of W. V. Quine*, (pp. 663–4). La Salle, Ill.: Open Court.
- Reichenbach, Hans (1938). *Experience and Prediction*. Chicago: University of Chicago Press.
- Schurz, Gerhard (2008). The Meta-Inductivist's Winning Strategy in the Prediction Game: A New Approach to Hume's Problem. *Philosophy of Science*, 75, 278–305.
- Wolpert, David H., & Macready, William G. (1997). No Free Lunch Theorems for Optimization. *IEEE Transactions on Evolutionary Computation*, 1, 67–82.