Kelly on Ockham's Razor and Truth-Finding Efficiency

Simon Fitzpatrick*†

This article discusses Kevin Kelly's recent attempt to justify Ockham's Razor in terms of truth-finding efficiency. It is argued that Kelly's justification fails to warrant confidence in the empirical content of theories recommended by Ockham's Razor. This is a significant problem if, as Kelly and many others believe, considerations of simplicity play a pervasive role in scientific reasoning, underlying even our best tested theories, for the proposal will fail to warrant the use of these theories in practical prediction.

1. Introduction. Ockham's Razor is the widely held scientific maxim that, other things being equal, we should prefer simpler theories to less simple ones. Simplicity, in this sense, is usually thought to be connected with positing fewer entities, causes, and laws, providing more unified explanations, and so forth. There appear to be many examples of Ockham's Razor at work in the history and current practice of science, Copernicus's appeal to the greater "simplicity" and "harmony" of heliocentric astronomy and the practice of curve fitting being some of the most commonly cited.

Despite its popularity, there is a long-standing philosophical puzzle surrounding the justification for Ockham's Razor. For many, a preference for simple theories can only make sense if we are entitled to assume that nature is simple, but since we do not seem to be so entitled, Ockham's Razor looks like mere wishful thinking. Although there have been valiant attempts (see Baker 2011), no fully satisfactory answer to this justificatory puzzle has been forthcoming.

Received December 2011; revised October 2012.

*To contact the author, please write to: Department of Philosophy, John Carroll University, 1 John Carroll Boulevard, University Heights, OH 44118; e-mail: sfitzpatrick@jcu.edu.

†Thanks to Earl Spurgin, Harry Gensler, and the referees for *Philosophy of Science* for helpful criticisms and suggestions. Research for this article was supported by a Summer Research Fellowship from John Carroll University.

Philosophy of Science, 80 (April 2013) pp. 298–309. 0031-8248/2013/8002-0003$10.00
Copyright 2013 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/670298 Published online by Cambridge University Press

However, in a series of recent papers, Kevin Kelly (2004, 2007, 2011) claims to have solved the puzzle. Using concepts from *formal learning theory*, Kelly argues that the methodological strategy of selecting the simplest theory compatible with the data can be shown to provide the most *efficient* method for finding the true theory, even if the truth is complex. "Efficiency" is to be understood in terms of minimizing theory changes en route to the true theory. Ockham's Razor thus makes sense, not because nature is simple but because it puts us on "the straightest possible path to the truth" (2007, 563).

Kelly's proposal is original and important but has not received the critical scrutiny it deserves. I will briefly review Kelly's efficiency argument and then describe what I take to be the most significant difficulty for his approach: the fact that it fails to warrant any confidence in the empirical content of theories recommended by Ockham's Razor. This is a significant problem, if, as Kelly and many others believe, considerations of simplicity play a pervasive role in scientific reasoning.

2. "The Straightest Possible Path to the Truth." Inspired by the work of Hilary Putnam and E. M. Gold, formal learning theory provides a mathematical framework for studying the complexity of empirical problems and for evaluating the effectiveness of particular inductive methods in solving them (see Osherson, Stob, and Weinstein 1986; Kelly 1996; Schulte 2012). Although most interest and work in this field has been confined to linguistics, mathematics, and computer science, Kelly and his collaborators, including Clark Glymour and Oliver Schulte, have sought to bring out the philosophical import of the approach, developing it into a general normative epistemological framework.[1] The key learning-theoretic concept that does the work in Kelly's proposal is *efficient convergence to the truth in the limit*.

A method of inductive inference, such as a rule of enumerative induction, is said to be convergent in the limit with respect to a particular empirical problem, such as the problem of determining whether all ravens are black, if and only if, given the background assumptions of the scientist, it is guaranteed to eventually settle on the true answer to the problem, no matter which of the possible answers consistent with background assumptions is in fact the true one. Kelly calls this *logical reliability* and takes it to be a basic desideratum for any inductive method, analogous to the requirement placed on computational algorithms that they are guaranteed to output the correct answer to a computational problem on every possible input. Of course, in the inductive case, an empirical hypothesis whose content outruns current data

1. Important contributions to the philosophical application of formal learning theory have also come from Daniel Osherson, Scott Weinstein, and their colleagues (see Osherson et al. 1986; Martin and Osherson 1998).

https://doi.org/10.1086/670298 Published online by Cambridge University Press

(e.g., a universal generalization) could always be refuted in the future. Thus, the requirement is only that an inductive method be guaranteed to eventually attain truth; we may never be sure when truth has been attained. Any such guarantee will hang on substantive background assumptions about the problem. For example, a rule of enumerative induction that, given observations of black ravens, conjectures "all ravens are black" until a nonblack raven is observed (whereupon it will conjecture "not all ravens are black") will be logically reliable for the problem of determining whether all ravens are black, so long as it is assumed that if there are any nonblack ravens they will eventually be observed.

However, for a given empirical problem there may be a potentially infinite number of methods that will eventually reach the truth. This is where considerations of efficiency come in. Computer scientists evaluate the efficiency of algorithms in terms of the number of steps taken before halting at the solution. The analogue for inductive inference in Kelly's framework is *retraction efficiency*: an inductive method is said to converge to the truth efficiently if, relative to other logically reliable methods, it minimizes the maximum number of retractions of answers (theory changes) that might be required of the scientist while using the method to guide her decisions, before settling on the true answer. Efficiency, in this sense, matters, according to Kelly, because scientists care about not just finding the truth but taking the least cognitively costly path to it—gratuitous theory changes being significant cognitive costs. Kelly's proposal is that, even if we do not know whether the truth is simple or complex, convergent methods that are biased toward simplicity can be shown a priori to be optimally efficient in this sense.

3. An Example. Nelson Goodman's (1983) New Riddle of Induction has been thought to illustrate the pervasive role of simplicity considerations in induction.[2] Given a sample of green emeralds, why do we prefer to conjecture that all emeralds are green rather than that all emeralds are grue, where something is grue if it is first observed before some future time t and is green or when it is first observed at or after t and is blue? Both hypotheses are equally consistent with observations of green emeralds before t, yet the grue hypothesis is not normally even considered. A common intuition is that the green hypothesis is simpler; hence, it is Ockham's Razor that underlies our natural inductive inclinations in this case. But, in what sense is the green hypothesis simpler than the grue one, and why should simplicity count?

For the moment, let us take for granted that the green hypothesis is simpler than the grue one. Suppose also that the world is such that either all

2. What follows is based on Schulte (1999, 2012). Readers should consult Kelly (2007, 2011) for refinements and extensions of the argument and to get a sense of the mathematics behind it.

https://doi.org/10.1086/670298 Published online by Cambridge University Press

emeralds are green or all emeralds are grue(t), where t is some time at and after which any newly examined emeralds are blue. Call *Ockham's Razor* the inductive method which says that, given observations of green emeralds, we should conjecture that all emeralds are green unless some blue emeralds are observed, in which case we should conjecture that all emeralds are grue(t), where t is the time at which blue emeralds are first observed. Compare this inductive method with a gruesome one, which conjectures that all emeralds are grue(2050), where the year 2050 is the specified value for t.

Both methods can be shown to be logically reliable. If all emeralds are green, Ockham's Razor will never abandon this conjecture. If all emeralds are grue(t), the method will switch to conjecture that at t. Hence, given our background assumptions, following Ockham's Razor will eventually get us to the true hypothesis, no matter what it is. The gruesome method will also be logically reliable, so long as it employs some rule for revising its conjecture if green emeralds continue to be found during or after 2050.

What about the relative efficiency of these methods? Ockham's Razor entails at most one retraction: the switch from "all emeralds are green" to "all emeralds are grue(t)," once a blue emerald is observed (if all emeralds are green, no retraction will be required). In contrast, the gruesome method may incur more than one retraction: for example, if green emeralds continue to be found during or after 2050, the method may switch to the conjecture that all emeralds are green, but it will still be possible that blue emeralds will be observed, whereupon a second retraction will be required. The same holds when we compare Ockham's Razor to other gruesome methods that specify different values for t. Of course, a scientist employing any of these methods might be lucky enough to live in world where the relevant grue(t) hypothesis is true and so will converge to the truth with fewer reversals than a scientist who follows Ockham's Razor. But, given that we do not know what the truth is at the outset, it is the worst case behavior of the methods that really matters, according to Kelly—thus, he makes no claim about the superiority of Ockham's Razor with respect to the expected number of retractions. Compared with all these gruesome methods, Ockham's Razor guarantees the lowest worst case costs in terms of retractions and lost time entertaining false theories. Hence, if one cares not just about finding the truth but finding it efficiently, one should follow Ockham's Razor.

A concern about appealing to Ockham's Razor here is that the relevant notion of simplicity might be an artifact of the language we use to express the theories. For example, the green hypothesis might be thought to be simpler because it posits that emeralds are "uniform" in their properties, whereas a gruesome hypothesis posits a lack of uniformity (some are green and some are blue). However, part of Goodman's motivation in posing the New Riddle was to show that such notions of uniformity are language variant. The green hypothesis can be redescribed using the predicates "grue" and "bleen,"

https://doi.org/10.1086/670298 Published online by Cambridge University Press

so that it posits a lack of uniformity in the properties of emeralds (some are grue and some are bleen).[3]

In response to this sort of concern, Kelly proposes that the complexity of a theory be defined relative to the structure of the empirical problem at hand: the ordering of simple to complex reflects the ordering of theories selected by a maximally efficient convergent method for solving the problem (for more precise definitions, see Kelly 2007, 2011). In this case, if the empirical problem is inferring the color of emeralds, then, given a sample of green emeralds, the green hypothesis comes out simpler than any gruesome one. Since this measure reflects the number of retractions that a scientist may be vulnerable to, it does not depend on the language that is used to express the theories. Thus, Kelly argues, "careful attention to truth finding efficiency provides not only a novel explanation of Ockham's Razor, but also a fresh perspective on the nature of simplicity itself" (2007, 570).

4. An Adequate Solution to the Simplicity Puzzle? The basic logic of Kelly's proposal is this: "Ockham's razor helps you find the truth without indicating what the truth is: disregarding Ockham's advice opens you to a needless, extra U-turn or reversal in opinion prior to all the reversals that even the best of methods would have to perform if the same answer were true. So you ought to heed Ockham's advice. Simplicity doesn't indicate the truth, but it minimizes reversals along the way. That's enough to explain the unique connection between simplicity and truth" (2004, 492).

This proposal has a number of attractions. It seems to shed light on the common intuition that, even if we do not know whether the truth is simple or complex, it is a good heuristic to start out simple and ratchet up complexity as new data require: starting out complex risks leading us down too many blind alleys before finding the truth. Kelly has shown that the analysis can be applied to the use of Ockham's Razor in various idealized empirical problems, including simple curve-fitting cases (Kelly 2007) and when isolating causes from correlation data (Kelly 2011). Impressive attempts have also been made to apply it to real cases from science—most notably, Schulte (2008) has applied it to the problem of discovering conservation laws in particle physics. The analysis has not been extended to statistical inference, where the relevant theories only make probabilistic claims about the data, but Kelly (2011) appears to be optimistic about doing so. The proposal also seems to avoid the wishful thinking, evasions, and circular arguments that plague so much of the existing literature on Ockham's Razor.

Nonetheless, important concerns can be raised about its adequacy. One is that the basic requirement of logical reliability seems highly restrictive:

3. Something is bleen if first observed before t and is blue or if first observed at or after t and is green.

https://doi.org/10.1086/670298 Published online by Cambridge University Press

how often is it that scientists occupy contexts in which their inferential methods can be guaranteed to get them to the truth given their background assumptions? This requirement also means that Kelly's proposal cannot justify choosing between empirically equivalent theories: in his framework no inductive method can be logically reliable in worlds where the data fail to rule out all but a unique theory in the limit. Scientific realists have sometimes invoked Ockham's Razor as a way to get around the purported problem of there being multiple incompatible theories consistent with all possible data. Kelly (2011, n. 13) is happy to reject such a realist view: in such worlds, all Ockham's Razor can do is get us to an empirically adequate theory. Realists, in contrast, may argue that logical reliability is too high a standard to set for the admissibility of inductive methods. Rather than pursuing these sorts of concerns, I want to focus on a more fundamental worry about Kelly's proposal, which concerns the nature of the justification he offers.

5. Retraction Efficiency: Epistemic or Pragmatic? A natural reaction to Kelly's proposal is that retraction efficiency appears to be a purely *pragmatic* consideration and not a genuinely *epistemic* one. How can concerns about possibly having to change one's mind a greater number of times provide good reason to believe, or even tentatively accept, the claims of the theory currently selected as the simplest one consistent with the data? Kelly is explicit that his efficiency argument solely concerns the worst case convergence properties of methods, not features of simpler theories themselves. It certainly does not show that simpler theories are more likely to be true. How, then, can simplicity be a genuine criterion for theory choice on Kelly's account?

Kelly argues that efficient convergence implies that scientists should adopt doxastic attitudes toward the theories recommended by their methods. A scientist who does not take the theories she conjectures seriously either will never converge to the truth because she never fixes belief or will be vulnerable to taking extra time deciding whether to change the epistemic attitude from mere conjecture to belief. Hence, the goal of efficient convergence provides a means-end justification for doxastic commitment to theories selected by Ockham's Razor: full-blooded Ockhamism is more efficient than half-hearted Ockhamism. Moreover, Kelly (2004, 2011) argues that this justification should be regarded as genuinely epistemic. For many epistemologists, what distinguishes epistemic from nonepistemic cognitive values is that their pursuit promotes the attainment of true beliefs. Hence, retraction efficiency should be regarded as a genuinely epistemic cognitive value: in contrast to paradigmatic pragmatic values, such as ease of understanding, its pursuit promotes speedy and stable attainment of true beliefs. Thus, even if simplicity cannot be regarded as a direct indicator of truth, we have a genuinely veristic justification for preferring simple theories.

https://doi.org/10.1086/670298 Published online by Cambridge University Press

Indeed, Kelly (1996, 2004) has sought to build a general account of epistemic justification out of the learning-theoretic notion of efficient convergence: *logical reliabilism*. Like other forms of reliabilism in epistemology, logical reliabilism analyzes the justification for a belief in terms of the reliability of the belief-forming process that led to its adoption. However, unlike traditional forms of reliabilism (e.g., Goldman 1979), reliability is not defined in terms of how often the process delivers true beliefs, rather than false ones, or in terms of it tracking the truth in the world the agent happens to live in. Instead, a belief is justified if it is selected by an efficient logically reliable method. Importantly, this account rejects the standard view in philosophy of science, according to which the justification for scientific theories is to be analyzed in terms of relations of confirmation between theories and empirical data. As a theory of scientific method, Kelly's approach is thus closer to the anti-inductivism of Popper since notions of confirmation, inductive support, and probability play no role. However, unlike Popperian falsificationism, it does sanction the adoption of doxastic attitudes toward scientific theories.

Critics should thus not beg the question by summarily dismissing Kelly's justification for Ockham's Razor as merely "pragmatic" and not genuinely "epistemic." Nonetheless, as I will now argue, serious concerns can be raised about the epistemic adequacy of this justification, which also cast doubt on the tenability of logical reliabilism more generally.

6. Strategy versus Confidence. In setting up the justificatory problem to be solved, Kelly endorses the common view in philosophy of science that simplicity considerations do not just play an occasional role in science—for example, in some explicit arguments given by scientists for preferring particular theories to others—but play a pervasive role in theory choice. Goodman's New Riddle, for example, has been thought to show that any scientific theory, no matter how well tested or established, can be confronted with a potentially infinite number of incompatible alternatives equally consistent with current data but which have different implications with respect to future experience—for example, theories that contain variants of the relevant laws that make different predictions for as yet unobserved times or places. Thus, according to Kelly, albeit often implicitly "Ockham's razor does the heavy lifting in scientific theory choice, for no other principle suffices to winnow the infinite range of possible explanations of the available data down to a unique one" (2011, 985). However, if Ockham's Razor does this much heavy lifting, significant concerns can be raised about the adequacy of Kelly's justification for it.

The worry is reminiscent of Wesley Salmon's (1981) criticism of Popperian falsificationism. Salmon emphasized the crucial role that scientific theories play in guiding practical action—the building of bridges, prescrib-

https://doi.org/10.1086/670298 Published online by Cambridge University Press

ing of medicines, and so forth. Both scientists and ordinary people talk and behave in ways that suggest our best tested theories furnish us with a confident (albeit not certain) basis for predicting future occurrences and, hence, provide a guide for how to best intervene in the world to achieve our immediate practical ends. Indeed, lives are frequently bet on the predictive accuracy of these theories—for example, every time an aircraft is judged flight worthy according to accepted principles of aerodynamics and aeronautics or when accepted nuclear theory is used to determine the safest place to build a nuclear reactor. Salmon's point was that Popper's account of the scientific method allowed that we may have preferences for certain theories over others, given the available data (those that are highly falsifiable and well "corroborated" by the data), but failed to warrant any confidence in the predictions of these best theories—Popper's notion of corroboration being explicitly only a measure of past predictive success. Hence, in cases in which immediate practical action is required, it could not explain why it is rational to follow the predictions of our best tested theories over those of any number of other unrefuted theories consistent with our experience to date.

Logical reliabilism faces the same problem. As we have seen, it allows that we may have preferences for certain theories over others, given the available data—those selected by efficient logically reliable methods. Unlike Popperian falsificationism, it also offers a means-end justification for doxastic commitment to these theories. However, this warrant for belief obtains only relative to the goal of efficient long-run convergence to the truth. Such long-run strategic considerations are completely unrelated to the short-run predictive accuracy of the theories in question and, thus, are not of the right sort to warrant the expectation that a chosen theory will do better than any other unrefuted theory in making immediate predictions about the future. It is not enough to argue that we might as well use the theories chosen by efficient logically reliable methods as guides for practical action since we are given no grounds for thinking that these methods will select theories that better serve our immediate practical ends than just choosing randomly among all the theories consistent with our experience.

Logical reliabilism thus cannot explain why it is rational to use our best theories in practical prediction, and it cannot underwrite the obvious confidence that scientists and ordinary folk have in the predictions of these theories.[4] Hence, insofar as it is meant to provide a general account of epi-

4. Kelly's framework can be applied no matter how the hypothesis space is restricted. If it is restricted such that all but one admissible theory can be ruled out by the data in the short run, we will have short-run grounds for confidence in the predictive content of the last theory standing, conditional on the relevant restricting assumptions. But it is surely rare for scientists to be in a position in which the hypothesis space can legitimately be

https://doi.org/10.1086/670298 Published online by Cambridge University Press

stemic justification in science without appealing to notions of confirmation or inductive support, logical reliabilism is deeply inadequate.[5] In addition to long-run convergence to the true theory, scientists have short-run goals, like practical prediction, which an account of theory justification must serve as well, but logical reliabilism leaves this out of the picture.[6] Moreover, if Ockham's Razor plays the fundamental role in theory choice that Kelly and many others think it does—such that it underlies preferences for even our best, most well-established theories—his strategic justification for this is insufficient. If it is ultimately simplicity that underwrites, for instance, a preference for current accepted theory in aerodynamics and aeronautics over some gruesome variant that predicts that aircraft judged flight worthy by standard theory will start falling out of the sky tomorrow, then an account of the justification for this preference that fails to warrant the (implicit or explicit) confidence that pilots and passengers flying tomorrow have in the predictions of standard theory should not seriously be seen as a solution to the problem of justifying Ockham's Razor.

7. Two Possible Replies. One reply that Kelly could make to these concerns is as follows. First, we simply do not have, and it seems unlikely that we will ever have, a solution to the problem of induction that can warrant confidence in what our theories say about unobserved events. Obviously, no finite sample guarantee can be had for the correctness of a universal generalization or other ampliative hypothesis. All that seems possible is some weaker notion of defeasible inductive support or confirmation, yet attempts to justify such notions have consistently failed to elude the Humean skeptic. This is partly why

restricted to this extent, and Kelly premises the importance of the problem of justifying Ockham's Razor on the notion that it is simplicity that underlies our preferences for our best scientific theories, not deduction from the data plus restrictions on the hypothesis space.

5. I noted earlier that Kelly's approach cannot justify some realist uses of Ockham's Razor. Here, logical reliabilism also appears inadequate from an antirealist perspective. Although contemporary antirealists deny we are entitled to be confident in the theoretical content of our best theories (i.e., what they say about unobservables), they usually hold that we are entitled to be confident in their empirical (observational) content, including what they predict about future observation (e.g., Van Fraassen 1980).

6. This is not a problem for more traditional reliabilist accounts of justification, which employ short-run notions of reliability. On Goldman's (1979) account, one is entitled to be confident in the predictive content of beliefs produced by reliable processes since for him a process is reliable if most of the beliefs it produces are true. Internalists about justification object that reliability is insufficient for justified belief since one may have no warrant for thinking that the process is reliable. However, internalist concerns do not motivate the current objection: Kelly's long-run notion of reliability fails to warrant confidence in the predictive content of a belief, whether or not we hold it necessary for justified belief that one has grounds for thinking that the belief-forming process is reliable.

https://doi.org/10.1086/670298 Published online by Cambridge University Press

Kelly (1996, 2004) prefers a different way of thinking about induction, in terms of procedures for finding the truth as efficiently as possible: we can realistically hope to pursue a priori investigation of which methods will facilitate efficient long-run convergence under various conditions. Second, attempting to connect simplicity more directly with truth looks like a fool's errand. For instance, Kelly (2011) shows that Bayesian accounts of simplicity claiming to link simplicity with high prior probability or likelihood all rely on question-begging assumptions. We should therefore learn to settle for truth-finding efficiency, and his proposal should be lauded for encouraging us to dismiss fantasies of having anything more.

We can agree that no fully adequate theory of inductive support or confirmation exists and that previous attempts to explicate a more direct link between simplicity and truth have been unsuccessful, but the question to ask is whether we can really accept as adequate an account of the scientific method and of Ockham's Razor that does not allow us to have any (defeasible) confidence in the predictive content of our best theories. Salmon's concern about Popper's attempt to excise confirmation from science was that "if the operation were successful, the patient (science), deprived of all predictive import, would die" (1981, 125). He did not deny the problem of induction. His point was rather that practical prediction is so central to scientific practice (and everyday life) that an account of the justification for the scientific method that fails to warrant any confidence in the predictive content of our best theories just would not be a solution to the justificatory problem. It would be like offering a "solution" to the problem of external world skepticism that provided no grounds for believing that one has hands. The same is true of logical reliabilism. It is not enough to claim that our best theories are warranted because choosing them will facilitate long-run convergence to the truth, if these theories play such a crucial role in guiding short-run practical action.[7] When it comes to Ockham's Razor, if simplicity really does play such a fundamental role in theory choice that it underlies even our best, most well-established theories, then the justification for this role, if one exists, has got to be of the sort that can motivate some level of confidence in the predictive content of these theories. Anything else is just not going to be fit for purpose.

Another quite different reply would be to suggest that logical reliabilism can motivate confidence in the empirical content of theories. Although Kelly rejects the idea that notions of confirmation and probability are fundamental to theory justification, he does allow that we may have a means-end ra-

7. We should also consider the use of our best theories as background theories, guiding the design and interpretation of experiments and the evaluation of other theories. This, scientific realists have often argued, implies confidence in their empirical and theoretical content.

https://doi.org/10.1086/670298 Published online by Cambridge University Press

tionale for employing such notions, if doing so can contribute to efficient convergence. For example, if prior probabilities can be appropriately distributed and the process of conditionalizing on the evidence implemented in such a way that a Bayesian agent will converge to the truth efficiently, a means-end justification will exist for that agent ascribing degrees of belief to theories and updating them in light of the evidence. Given this, a probabilistic measure of confidence in the empirical content of theories could perhaps be provided, underwritten by considerations of truth-finding efficiency.[8]

This will not help. The warrant for such a confidence measure will still have nothing to do with the short-run predictive accuracy of the theories in question. We will have a rationale for ascribing particular degrees of belief to theories relative to the goal of efficient long-run convergence to the truth but not relative to the goal of effective short-run practical prediction. Once again, we will be left with no grounds for thinking that our best theories will better serve our immediate practical ends than any of the other unrefuted theories consistent with our experience.

8. Conclusion. The preceding discussion suggests that the simplicity puzzle may be deeper than is commonly recognized. On the one hand, it seems extremely difficult to show how simplicity could be correlated with truth or even just empirical adequacy. On the other, if Ockham's Razor is as important to the scientific method as Kelly and others have suggested, we can settle for nothing less.

One possible way out is to deny that simplicity is so fundamental. Elliott Sober (1988) has argued that in many cases in which simplicity seems to be important in theory choice, something else—like local background knowledge—does the real epistemic work. Such a deflationary view of simplicity is at least worth exploring further, if it can relieve us of the task of establishing a general correlation between simplicity and truth.

In any event, truth-finding efficiency may still help us make sense of some uses of Ockham's Razor in science. Kelly's work does advance our understanding of the epistemic benefits a bias toward simplicity may bring. He has shown that there is something to the idea that preferring simple theories is a rational strategy, and this is not for purely pragmatic reasons—at least,

8. Formal learning theorists have been highly critical of Bayesianism (Juhl 1993; Kelly 1996; Martin and Osherson 1998), arguing, e.g., that Bayesian norms of rationality sometimes inhibit logical reliability. However, this does not mean that Bayesian methods cannot be put to the service of efficient convergence, under appropriate conditions. It remains an open question what those conditions are, but Kelly (2011) provides some preliminary remarks on how the efficiency argument could be applied to Bayesian methods that output degrees of belief.

https://doi.org/10.1086/670298 Published online by Cambridge University Press

not for reasons entirely unconnected with truth. My concern here, however, has been to show that if, as Kelly and many others believe, simplicity plays a pervasive role in scientific theory choice, truth-finding efficiency fails to provide the right justification for this role.

REFERENCES

Baker, Alan. 2011. "Simplicity." In *Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. Stanford, CA: Stanford University. http://plato.stanford.edu/archives/sum2011/entries/simplicity/.

Goldman, Alvin. 1979. "What Is Justified Belief?" In *Justification and Knowledge*, ed. George S. Pappas, 1–23. Dordrecht: Reidel.

Goodman, Nelson. 1983. *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press.

Juhl, Cory. 1993. "Bayesianism and Reliable Scientific Inquiry." *Philosophy of Science* 60:302–19.

Kelly, Kevin. 1996. *The Logic of Reliable Inquiry*. New York: Oxford University Press.

———. 2004. "Justification as Truth-Finding Efficiency: How Ockham's Razor Works." *Minds and Machines* 14:485–505.

———. 2007. "A New Solution to the Puzzle of Simplicity." *Philosophy of Science* 74 (Proceedings): 561–73.

———. 2011. "Simplicity, Truth, and Probability." In *Handbook on the Philosophy of Statistics*, ed. Prasanta Bandyopadhyay and Malcolm Forster, 983–1024. Oxford: North-Holland.

Martin, Eric, and Daniel Osherson. 1998. *Elements of Scientific Inquiry*. Cambridge, MA: MIT Press.

Osherson, Daniel, Michael Stob, and Scott Weinstein. 1986. *Systems That Learn*. Cambridge, MA: MIT Press.

Salmon, Wesley. 1981. "Rational Prediction." *British Journal for the Philosophy of Science* 32: 115–25.

Schulte, Oliver. 1999. "Means-Ends Epistemology." *British Journal for the Philosophy of Science* 50: 1–31.

———. 2008. "The Co-discovery of Conservation Laws and Particle Families." *Studies in History and Philosophy of Modern Physics* 39:288–314.

———. 2012. "Formal Learning Theory." In *Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. Stanford, CA: Stanford University. http://plato.stanford.edu/archives/sum2012/entries/learning-formal/.

Sober, Elliott. 1988. *Reconstructing the Past: Parsimony, Evolution, and Inference*. Cambridge, MA: MIT Press.

Van Fraassen, Bas. 1980. *The Scientific Image*. Oxford: Oxford University Press.

https://doi.org/10.1086/670298 Published online by Cambridge University Press