

THE RELIABILITY CHALLENGE AND THE EPISTEMOLOGY OF LOGIC

Joshua Schechter
Brown University

1. Introduction

This paper concerns a problem in the epistemology of logic. This problem is an analogue of the Benacerraf-Field problem for mathematical Platonism. It is also an analogue of Mackie's queerness argument against moral realism, at least on one way of understanding that argument. The problem is what I call the "Reliability Challenge" for logic.

The purpose of this paper is not to solve this problem – although I do gesture toward a solution. Rather, my aim is to get clear on what exactly the problem is. It is my view that the nature of the reliability challenge has not been properly understood. My main goals here are to present the reliability challenge in its strongest possible form and to show why two seemingly attractive responses are unable to answer it. This will prove to be relevant not only to the epistemology of logic, but also to the epistemologies of mathematics, modality, morality, and other *a priori* domains.

Let me begin by stating a few claims about logic. Certain propositions are logical truths. Such truths include the proposition that every walrus is a walrus and the proposition that if both it is raining and if it is raining then the roads will be slippery, then the roads will be slippery. Other propositions are logical falsehoods. Such falsehoods include the proposition that some walrus is not a walrus and the proposition that it is raining and if it is raining then the roads will be slippery, but the roads will not be slippery.

For my purposes here, I need not assume any particular account of the nature of logic. In my discussion, however, I will make use of a few background assumptions. First, “logic,” as I use the term here, does not concern an artificial formal language but propositions that can be expressed in natural language and believed by ordinary thinkers.¹ Second, I assume that propositions (and not sentences) are the primary bearers of logical truth and logical falsity.² Third, I assume that logical truths are necessarily true and logical falsehoods are necessarily false on any reasonable (alethic) kind of necessity. Finally, I assume that classical logic is at least approximately correct.

We ordinarily think that we know many logical truths. For instance, we take it that we know the two logical truths stated above. At first blush, this knowledge is unmysterious. What could be easier than knowing that every walrus is a walrus? Yet, we ought to be more puzzled than we are. Our logical knowledge is a significant cognitive achievement in at least two respects: First, we are *reliable* in by-and-large believing logical truths and disbelieving logical falsehoods. Second, our logical beliefs have a positive epistemic status. We are justified – or perhaps better, *epistemically responsible* – in holding many of the logical beliefs that we do. That we are reliable and that we are responsible are facts that need to be explained. It would be very odd to think that either is a brute fact, akin to a law of nature or to a mere random accident. But it is not at all clear how these facts are to be explained.

The epistemology of logic thus has two main explanatory tasks – to explain how it is that our logical beliefs are reliable and to explain how it is that we are epistemically responsible in believing as we do. In this paper, I focus on the first task, that of explaining our reliability.³

For simplicity, let the term “logical propositions” stand for the logical truths and logical falsehoods. We can understand the claim that we are reliable about logic as the following thesis: The logical propositions we believe (upon reflection and discussion) are by-and-large true and the logical propositions we disbelieve (upon reflection and discussion) are by-and-large false.⁴ The reliability challenge for logic is the challenge of explaining this fact.⁵

Reliability challenges are not limited to the case of logic. Given any domain about which we think we are reliable, there is a need to explain our reliability.⁶ For certain domains – for instance, facts concerning ordinary-sized objects in our environment – we are able to answer the challenge. For other domains, the challenge remains pressing. Significant reliability challenges arise for mathematics, morality, modality, conceptual truths, intuitions about thought experiments, and many other domains.

The intuitive problem can be put as follows: We have some understanding of how perception can yield veridical beliefs about the external world. We possess a sketch of how the mechanisms underlying perception work and understand how they may yield true beliefs. But this explanation doesn’t extend to the cases of logic, mathematics, modality, and other *a priori* domains.⁷ Nor is there available any well-developed alternative account. We simply do not understand how we can be reliable about these domains, given that our beliefs were not arrived at via some kind of perception.

Reliability challenges are perhaps most familiar for the cases of mathematics and morality. According to the Benacerraf-Field problem, mathematical Platonism faces difficulty in explaining the reliability of our mathematical beliefs, given its claim that mathematical entities are abstract.⁸ According to one version of Mackie’s queerness

argument, moral realism faces difficulty in accounting for our reliability about moral truths, since if objective values were to exist, they would be peculiar entities outside of our ken.⁹

These difficulties are typically presented as generated by the ontologies of the two domains. But this is a mistake. The root of the trouble is not the ontology but the apparent *objectivity* of mathematics and morality.¹⁰ If mathematics, for instance, were to turn out not to have an ontology, but the relevant truths were nevertheless objective, our reliability would remain puzzling. If mathematics were to turn out to have an ontology, but the relevant truths were not objective – for instance, they were somehow constituted by our practices – our reliability might be easily explained. The same holds true for other domains. For the case of logic, what gives the reliability challenge its bite is the apparent objectivity of logic.

There may be several philosophically interesting notions of objectivity.¹¹ For the purpose of developing the reliability challenge for logic, the claim that logic is objective can best be understood as the conjunction of three theses. The first thesis is a claim about the content of our sentences and mental representations:

Meaningfulness: Certain sentences and mental representations express logical truths and falsehoods. They are therefore both meaningful and truth-apt.¹²

This excludes nonsense, make-believe, and mere expressions of emotion as counting as objective. It provides an extremely weak requirement on the objectivity of logic.

The second thesis is a claim about independence:

Independence: The truth of logical truths and the falsity of logical falsehoods do not depend on us. In particular, they do not depend on our thoughts, language, or social practices.

For my purposes here, it suffices to operate with a purely intuitive understanding of this thesis. In my view, however, the best way to understand the independence thesis is as follows: The logical facts do not obtain (in whole or in part) in virtue of facts about us. Here, the *in virtue of* relation is an explanatory relation. Our thoughts, language, and practices are not part of what constitutively explains why the logical facts are the way that they are.¹³

The third and final thesis is the denial of a plenitudinous view of logic:

No Plenitude: Of the many possible different logical practices, only a few are correct.

This thesis ensures that reliability is not a trivial accomplishment – not just any logical practice counts as correct. Indeed, most logical practices do not.¹⁴

To motivate the need for the no-plenitude thesis, it may help to consider the cases of mathematics and morality. One answer that has been proposed to the Benacerraf-Field problem is roughly as follows: The domain of mathematical reality is so large and multifarious that no matter which consistent mathematical theories we were to accept, they would truly represent some portion of mathematical reality.¹⁵ This view satisfies the first two objectivity theses for mathematics. Yet, it is intuitively a non-objective view.¹⁶ Similarly, consider an analogous view about morality. Suppose that there are a vast number of rightness properties, all on a par, so that any consistent theory of right and wrong will truly characterize one of them. Surely, such a view of morality should not count as an objective view.

That this conception of objectivity is the appropriate one for the purpose of developing the reliability challenge may be motivated as follows: Suppose we are reliable about some *a priori* domain. There would seem to be only a few candidate explanations

of our reliability: (i) Our beliefs depend on the relevant facts; (ii) The facts depend on our beliefs; (iii) Both our beliefs and the facts depend on some third factor; and (iv) Reliability about the domain is trivial in that we would be reliable no matter which practice we were to adopt.¹⁷ The *a priori* of the domain would seem to rule out the first candidate. The independence condition would seem to rule out the second. The *a priori* and independence conditions together would seem to rule out plausible versions of the third. The no-plenitude condition would seem to rule out the final candidate. This is what generates the puzzle.

As we will see, this is an overly simplistic way to think about the challenge. But it provides a useful heuristic.

The difficulty of answering the reliability challenge for logic potentially has significant ramifications. If we were to come to believe that there is no satisfying explanation of our reliability compatible with the objectivity of logic and with our general background views about the world, this would put pressure on our belief that logic is objective, on our belief that we are reliable, or on our general background views.¹⁸ But giving up any of these claims would be devastating to our ordinary ways of thinking. Thus, there is an important explanatory challenge to answer.

This paper will proceed as follows: In the next section, I further clarify the reliability challenge by motivating a shift in focus from the reliability of our beliefs to the reliability of our cognitive mechanisms. In section three, I consider an argument due to Lewis that suggests that reliability challenges are mere pseudo-problems. This yields a significant refinement of the challenge. Section four is devoted to considering and rejecting several lines of thought that suggest that an answer to the reliability challenge is

unnecessary. In section five, I evaluate two potential responses to the reliability challenge – one based on rational insight and one based on the nature of concepts and concept-possession. I also discuss the prospects of an evolutionary account. Finally, in section six, I discuss how the case of logic may differ from that of mathematics, modality, morality, and other *a priori* domains.

2. The Reliability Challenge for Logic

Let us try to get the reliability challenge for logic in better focus. As a first pass, the challenge can be understood as that of answering the following question:

The Reliability Challenge (1): What explains our reliability about logic? That is, what explains the fact that the logical propositions we believe (upon reflection and discussion) are by-and-large true and the logical propositions we disbelieve (upon reflection and discussion) are by-and-large false?

There is a straightforward answer to this question: We are reliable in our logical beliefs because we employ reliable deductive rules of inference. Let me explain.

In order to understand how we are reliable in our logical beliefs, it is useful to first examine how we come to have such beliefs. Consider some moderately complex logical truth, for instance, *if both A and B then either B or C*, substituting particular sentences for *A*, *B*, and *C*. We believe this proposition, at least upon reflection. How did we come to believe it? In some cases, thinkers may come to believe this truth on the basis of testimony – perhaps of a logic teacher or a guru. But the canonical way in which thinkers come to believe this truth is via a chain of reasoning, perhaps one such as the following:

Suppose both *A* and *B*.
So *B*.
So either *B* or *C*.

So if both A and B then either B or C.

In general, we believe logical truths and disbelieve logical falsehoods on the basis of deductive reasoning. Our logical beliefs are the outputs of deductive reasoning in cases where there are no initial premises.¹⁹

What is going on when someone reasons in this way? On a simple picture, reasoning is the process of adding and subtracting to one's stock of beliefs. But reasoning is in fact much more complex than this simple picture suggests. Reasoning may involve suppositions, such as in the deduction displayed above. It may involve updating the collection of propositions a thinker rejects, or the degree to which a proposition is accepted. Moreover, a thinker may engage in a line of reasoning even if the thinker already believes claims made along the way. So what may change in an inferential step is not what is believed, but the basis on which a belief is held.

No matter how complex reasoning may be, it is plausible that whenever we rationally update our mental states, we do so via the employment of rules.²⁰ There are two principal grounds for this claim. First, appealing to rules is the most promising strategy for explaining the difference between genuine reasoning and mere change in belief. Second, we are familiar with two different sorts of mistakes thinkers may make in their reasoning – errors of competence and errors of performance. Appealing to rules helps to explicate this distinction; thinkers may employ the wrong rules, or they may misapply the rules that they employ.

Deductive reasoning is reasoning that involves deductive rules of inference. These are rules that are intimately connected to the logical concepts. It is an interesting question which deductive rules of inference ordinary thinkers employ. The simple deduction above

fits with the view that the rules we employ resemble the rules that appear in Natural Deduction formulations of logic. Indeed, such formulations were developed, in part, to connect logic more closely with psychologically realistic patterns of inference.²¹ On this view, the deductive rules we employ include rules such as:

From *both p and q*, infer *p*.
From *both p and q*, infer *q*.
From *p*, infer *either p or q*.
From *q*, infer *either p or q*.
From *q* under the supposition *p*, infer *if p then q*.

There are reasons to think that the deductive rules of inference we employ are much more complex than this list suggests.²² Nevertheless, it is appealing to think that we employ deductive rules that are closely related to the standard Natural Deduction rules, and in what follows I will assume that we do.²³

What, then, explains the reliability of our logical beliefs? Consider again *if both A and B then either B or C*. We believe this truth because we went through something like the chain of reasoning displayed above, and the transitions involved in this reasoning yield truths. The transitions yield truths because the deductive rules that govern them are reliable. In general, we are reliable in our logical beliefs because we are reliable in our deductive reasoning. Our reliability about logic is a byproduct of a more basic deductive competence.²⁴

This explanation answers the question that we started out with, Reliability Challenge (1). But it does not fully answer the reliability challenge for logic. It merely pushes the explanatory demand one step back. An explanation is now required of the reliability of our deductive reasoning. The challenge becomes that of answering the following question:

*The Reliability Challenge (2): What explains the reliability of our cognitive mechanism for deductive inference?*²⁵

This is a more fundamental – and much more general – explanatory challenge. It concerns not only the deductive reasoning that yields logical beliefs, but deductive reasoning from arbitrary *a priori* and *a posteriori* premises, too.

Reliability challenges for other domains are similar. The reliability of our mathematical beliefs, for example, can be explained by appeal to the reliability of the mechanisms underlying mathematical reasoning. The reliability of our moral beliefs can be explained by appeal to the reliability of the mechanisms underlying moral reasoning. In general, what primarily requires explanation is not the reliability of our beliefs, but the reliability of the cognitive mechanisms that generate them.

3. Lewis's Response

To further clarify the nature of the reliability challenge for logic, it is helpful to consider the radical view that the challenge is a mere pseudo-problem. Such a view is suggested by Lewis's discussion of mathematical and modal knowledge in *On the Plurality of Worlds*.²⁶ In that work, Lewis argues that there is only a challenge in explaining our reliability about contingent truths. For contingent truths, an account is needed of how we track the relevant truths. For necessary truths, in contrast, no such account is needed.

Simply believing the relevant truths is an infallible way to get it right. As Lewis writes:

[I]f it is a necessary truth that so-and-so, then believing that so-and-so is an infallible method of being right. If what I believe is a necessary truth, then there is no possibility of being wrong. That is so whatever the subject matter of the necessary truth and no matter how it came to be believed. (Lewis, 1986, 114–115)

Any residual feeling that there is something left to explain is presumably the result of drawing an illegitimate analogy between the necessary and the contingent.²⁷

Applying this line of thought to the case of logic, the idea is that since logical truths are necessary truths, there is no need to explain our reliability about logic. Simply believing the logical truths is an infallible way to get it right.

There is something attractive about Lewis's suggestion. There does seem to be a significant difference between the epistemology of the necessary and the epistemology of the contingent. But it is difficult to swallow the claim that there is no reliability challenge for necessary domains. Consider the case of a thinker who reliably believes some very complex logical or mathematical truths but who does not otherwise evidence any mathematical acumen. Surely, some explanation is needed of how this thinker is reliable. Similarly, consider familiar *a posteriori* necessary truths, such as the claim that water is molecular. Our coming to believe such claims was a substantial achievement, and so, surely, an account is needed of how it is that we got them right.

It is widely accepted that Lewis is mistaken, and that there is a reliability challenge for logic, mathematics, and other necessary domains.²⁸ However, no one has yet identified the central problem with Lewis's discussion. Identifying this problem will help us to clarify the nature of the reliability challenge.

A preliminary point to make in response to Lewis is that reliability challenges do not apply, in the first instance, to beliefs. As we have seen, the reliability of our beliefs can be explained in terms of the reliability of the underlying cognitive mechanisms that generate them. The challenge, rather, is to explain the reliability of these mechanisms.

What is it to explain the reliability of a cognitive mechanism? There is an important distinction that ought to be drawn here, one that seems to have been largely overlooked. Namely, there are two distinct explanatory demands concerning the reliability of a cognitive mechanism. This contrast is perhaps easiest to illustrate for the case of visual perception.

The first demand is to explain how our visual mechanism works such that it is reliable. A sketch of the answer to this question is known. There is a complex psychophysical story that explains the workings of our visual system. It explains how certain inputs – stimulations of our optic nerve – lead to the production of certain outputs – representations of the world. Another complex physical story explains how photons interact with objects and then come to stimulate our optic nerve. When these stories are put together, they entail that beliefs formed on the basis of visual perception are reliable indicators of the state of the external world, at least in environments like ours. This explains why our visual mechanism is reliable.

This explanation, however, does not explain all there is to explain. There is also the question of how it is that we have a reliable visual mechanism. Here, too, a sketch of the answer is known: We have a reliable visual mechanism because, very roughly, it conferred a heritable survival or reproductive advantage on our ancestors to correctly represent their environment using vision.

In general, there are two distinct questions concerning the reliability of any cognitive mechanism. For the case of deductive inference, they can be put as follows:

The Operational Question: How does our cognitive mechanism for deductive inference work such that it is reliable?

The Etiological Question: How is it that we have a cognitive mechanism for deductive inference that is reliable?²⁹

These two questions ought not to be identified with the question of how the relevant cognitive mechanism works and the question of how we came to have it. Explaining how a mechanism works or how we came to have it is neither necessary nor sufficient to explain its reliability in either of the two senses.

To fully answer the reliability challenge for logic, satisfying answers to both the operational and etiological questions must be provided. At a minimum, what's required are sketches of plausible answers. An inability to sketch a plausible answer to one of the two questions would put pressure on the claim that our deductive mechanism is reliable, the claim that logic is objective, or on our general background views of the world. And if we had reason to think that a satisfying answer was impossible, our epistemic situation would be worse still.

Let's examine these two questions in turn. Consider first the operational question. To answer the analogous question for the case of vision, a complex causal story must be provided, one that explains how we track the relevant empirical facts. By analogy, one might think that a causal tracking story must be provided for the case of deductive inference, too. But this would be a mistake. A simpler kind of explanation suffices. Our deductive mechanism works via the employment of deductive rules of inference. The mechanism is reliable because the deductive rules we employ are necessarily truth-preserving. That's all that needs to be said; no causal interaction with the environment or with some mysterious realm of logical facts need be invoked.³⁰

Now consider the etiological question. This question is not subject to so simple a response. There was certainly no guarantee that we would come to have reliable

deductive rules. To answer this question, a non-trivial account is required. This is the crux of the reliability challenge for logic.

The cases of mathematics, modality, and morality are similar. For each of these domains there is a straightforward way to answer the operational question: Our cognitive mechanisms are reliable because they involve the employment of necessarily truth-preserving rules. The pressing challenge is that of answering the etiological question.

Indeed, this challenge generalizes further still. Etiological challenges can be raised for the cognitive mechanisms involved in empirical reasoning, perhaps most notably, the cognitive mechanisms governing our inductive practices and the cognitive mechanisms governing our applications of concepts. For each of these cognitive tasks, there is the question of how it is that we possess reliable mechanisms.

We are now in a better position to evaluate Lewis's discussion. Lewis is correct in his claim that the reliability challenge for mathematics, modality, and related domains is subject to a straightforward response, so long as the challenge is construed to be that of answering the operational question.³¹ Lewis is mistaken in his claim that a non-trivial account is unneeded. There remains the challenge of answering the etiological question for each of these domains. In particular, for the case of logic, it must be explained how it is that we employ reliable deductive rules of inference.

4. Is an Explanation Really Needed?

Now that we have a better fix on the nature of the reliability challenge for logic, it is time to examine ways in which this explanatory demand might be resisted. There are several

lines of thought that suggest that an explanation of our reliability is unneeded or is easy to provide. The purpose of this section is to show that these lines of thought are mistaken.

One might think that an answer to the reliability challenge is unneeded. After all, thinkers can gain knowledge of their surroundings via visual perception despite being wholly ignorant of how vision works or how they came to have a reliable visual faculty. Indeed, ancient peoples had no real understanding of vision but nevertheless were able to rely upon it to gain knowledge about the world. Thus, it might be concluded that the cost of failing to answer a reliability challenge is extremely low.

In general, failing to possess an explanation of the reliability of a cognitive mechanism is not terribly worrisome. What would be worrisome is if we had compelling reason to suspect that there is no satisfying explanation to be had. To illustrate this, consider Field's example of someone who claims that their beliefs about the daily events in a remote village in Nepal are by-and-large true.³² Field argues that if we had good reason to doubt that there is a mechanism by which this could be the case, we should be very suspicious of the claim. There is a natural second condition: If we had good reason to doubt that there is a satisfying explanation of how it is the thinker came to have such a reliable mechanism, we should also be suspicious of the claim. In general, reason to doubt that there is a satisfying answer to the operational or etiological question for an (apparently) objective domain puts pressure on our overall view of the world.

At least *prima facie*, there is compelling reason to doubt that there is a satisfying answer to the etiological question for the cases of logic, mathematics, morality, modality, and other *a priori* domains. Each of these domains is plausibly objective. We seem to be reliable about each of these domains. But it is difficult to envision satisfying accounts of

how it is that we have reliable cognitive mechanisms for reasoning about these domains. Explanations of how it is that we have reliable cognitive mechanisms for reasoning about these domains – or at the very least, sketches of such explanations – are thus sorely needed.

There is a second line of thought that suggests that an explanation of our reliability about logic may be unneeded. Not every fact requires explanation. There are, after all, brute facts. Why couldn't we simply claim that it is a brute fact that we have a reliable cognitive mechanism for deductive inference? Wouldn't this be an acceptable answer to the etiological question for logic?

Brute facts come in two general kinds – fundamental laws and mere accidents. The former class plausibly includes laws of physics and metaphysics (if such there be). The latter class plausibly includes the initial conditions of the universe as well as random phenomena such as the exact timing of events of radioactive decay.

Our possession of a reliable deductive mechanism is not a plausible candidate for being a fundamental law. It would be strange to claim that it is a fundamental law that certain of our mental processes are truth-conducive. We don't think that there are fundamental laws in psychology or that such laws can require that our reasoning be reliable.

Our possession of a reliable deductive mechanism is also not a plausible candidate for being merely an accident. The reason is that our possession of a reliable deductive mechanism is a striking fact.³³ It "calls out" for explanation. *Ceteris paribus*, it is a cost of a theory if it treats striking phenomena (within the domain of the theory) as merely

accidental or otherwise unexplained. So it is a significant cost of a view if our possession of a reliable cognitive mechanism is left unexplained.

There is a further reason why it would be unsatisfying to claim that our possession of a reliable deductive mechanism is merely an accident. If it was accidental that we came to have a reliable deductive mechanism, it was presumably highly unlikely for us to have ended up with a reliable mechanism; there are far more ways to be unreliable than to be reliable. Accepting that our reliability came about by accident would therefore put pressure on our overall view of the world.

A third line of thought suggests that the reliability challenge for logic can, at least in principle, be easily answered. Presumably, there is some causal story that explains why we employ the particular deductive rules of inference that we do. It is a necessary feature of these rules that they are truth-preserving. So, the thought goes, the causal story will explain why it is that we employ reliable deductive rules. It will answer the etiological question for logic.³⁴

The trouble with this line of thought is that explanation is not closed under necessary consequence – or even logical entailment. For instance, suppose it were true that every day in March, the number of people who took the New York subway was a prime number. Suppose we could provide an elaborate explanation of why various people did (or did not) take the subway on particular days in March. Even though this explanation would entail that a prime number of people took the subway each day in March, it would not explain this fact. The fact that there was always a prime number of people would remain mysterious.³⁵

Analogously, an explanation of why it is that we employ particular deductive rules would not *per se* explain why it is that we employ reliable deductive rules. The explanation of our employment of the particular rules would presumably have nothing to do with their reliability. So the fact that we employ reliable rules would remain mysterious.

There is a final line of thought worth discussing. The idea is that deductive reasoning is epistemically special in that there is simply no room for the doubt that our deductive mechanism is unreliable. Thus, the reliability challenge could not undermine our logical beliefs. So there is no need to answer the reliability challenge for the particular case of logic.

There are several grounds for the claim that deductive reasoning is special in this way. The first is that there is reason to think that we cannot rationally have less than full confidence in our fundamental rules of reasoning. We could rationally have less than full confidence in some fundamental rule only if our fundamental rules recommended that we not fully trust it (given the appropriate input beliefs). In such a situation, our epistemic system would make two incompatible recommendations: It would (implicitly) recommend employing the rule and it would (explicitly) recommend not fully trusting it. It is incoherent for a thinker's epistemic system to issue such incompatible recommendations. Therefore, this line of thought goes, it is impossible for a thinker to mistrust one of their fundamental rules. Insofar as our deductive rules of inference are among our fundamental rules of reasoning, there is no room for doubting their reliability.³⁶

The second ground for thinking that deductive reasoning is epistemically special stems from the thought that any argument against the reliability of our deductive rules would have to rely (at some stage) on a deductive inference. Such an argument would therefore be self-undermining. It would be incapable of being rationally persuasive.

The third and final ground is simply this: It is difficult to take seriously the claim that a philosophical argument could challenge the reliability of a simple deductive inference, say the inference from a conjunction to one of its conjuncts. The reliability of such an inference is much more certain than any philosophical argument could be.³⁷

Perhaps deductive reasoning *is* special in this way. Perhaps the reliability of our deductive mechanisms cannot seriously be in question.³⁸ Even supposing that this is so, it does not show that an answer to the reliability challenge for logic is unneeded. Even if there is no room to doubt our reliability, that we are reliable is a striking fact. It calls out for explanation. Any reason to suspect that there isn't an explanation generates a tension in our overall view of the world. There is pressure to alleviate this tension by doing at least one of the following things: (i) giving up the claim that we are reliable; (ii) giving up the claim that logic is objective; or (iii) giving up some of our general background views. Even if it is rationally impossible to reject the claim that we are reliable, there remains the possibility of rejecting the objectivity claim or some of our background views. And even if there is not all-things-considered reason to give up any of these views, in the absence of a response to the reliability challenge, an uncomfortable tension in our overall view of the world would remain.

5. Candidate Explanations

The reliability challenge for logic is pressing. How might it be answered? So far as I can tell, there are only three proposals that are initially plausible – one involving a faculty of rational insight, one involving the nature of concepts and concept-possession, and one involving evolution. Let us examine each of these proposals in turn.³⁹

Rational Insight

One view is that our reliability about deductive inference can be explained by appeal to a kind of rational insight. The idea is that we have a cognitive faculty broadly analogous to perception by which we “see” that certain deductive inferences are good, or that certain propositions ought to be accepted, or something similar.⁴⁰

There are two familiar criticisms of rational insight-based views.⁴¹ First, there is no independent evidence that we have a faculty of rational insight. While it is plausible that we have mental states such as intuitions or “intellectual seemings,” these are best understood as felt inclinations to believe rather than the products of a distinctive cognitive faculty.⁴² Second, the mechanism behind rational insight is very mysterious. If rational insight is understood on the model of vision, our logical beliefs or deductive inferences would have to be *caused* by whatever it is that we have rational insight into. No plausible causal story of this sort has ever been developed. And no one has come up with a coherent alternative model.

But there is a still more serious problem with an appeal to rational insight, at least in the current context. Namely, appeal to a faculty of rational insight is of *no help whatsoever* in answering the etiological question for deductive inference. As I argued above, there is a straightforward explanation of how our deductive mechanism works

such that it is reliable. What requires explanation, rather, is how it is that we have a reliable mechanism. Appealing to a faculty of rational insight does nothing to meet this challenge. If anything, it increases its difficulty, since an explanation is then owed of how it is that we have a reliable faculty of rational insight. No such explanation seems forthcoming.

Appealing to a faculty of rational insight thus seems impotent to answer the reliability challenge for any domain.

Concept-Constitution

A second proposal is that our deductive reliability can be explained purely by virtue of the nature of concepts and concept-possession.⁴³ This suggestion involves two central ideas.

The first idea is that the deductive rules we employ help to constitute the meanings of the logical concepts we possess in the following sense: to fully master these concepts, a thinker must employ the rules.⁴⁴ For instance, to possess the concept of conjunction, thinkers must employ the rule: from *both p and q*, infer *q*. This view is thus committed to a version of Conceptual Role Semantics for the logical concepts.⁴⁵

The second idea is that the semantic values of logical concepts are assigned in such a way as to ensure that their constitutive rules are necessarily truth-preserving. In other words, the meaning of a logical concept is assigned so as to make its constitutive rules reliable.

Taking these two ideas together, we have a straightforward explanation of how it is that we employ reliable deductive rules: The deductive rules we employ are concept-constituting, and thus are guaranteed to be reliable.

More generally, the strategy of explanation is as follows: Certain rules we employ are constitutive of genuine concepts that we possess. The semantic values of genuine concepts are assigned in such a way as to make their constitutive rules reliable. This explains how it is that we employ reliable rules.

This view deserves serious consideration. But there is a significant problem with it, one that shows that it cannot fully answer the reliability challenge.

We can approach this problem by noting two general facts about concepts, assuming that some version of Conceptual Role Semantics is true. The first fact is that not every conceptual role – that is, not every package of rules involving a term in the mental lexicon – is constitutive of a genuine concept.

For an extreme case, consider Prior's example of "tonk".⁴⁶ The conceptual role governing this term is as follows:

From p , infer p tonk q .
From p tonk q , infer q .

It is very plausible that this conceptual role does not constitute a genuine concept. We do not understand what the word "tonk" means. ("Tonk" can only sensibly be mentioned and not used.) We do not understand what it would be to have tonk-ish thoughts.

For "tonk", there are several candidate explanations of why it does not stand for a genuine concept – for instance, it trivializes any conceptual practice to which it is added. But there are also other, less extreme examples of the same phenomenon that are more difficult to explain. Consider the conceptual role of any logical concept modified by

adding some *ad hoc* restriction.⁴⁷ For instance, consider the standard conceptual role for conjunction with the added restriction that one of its elimination rules (say, from *p and q*, infer *q*) only applies to sentences with a prime number of singular terms. This modified bundle of rules does not correspond to a genuine concept – we have no idea how to understand the corresponding term. Indeed, most conceptual roles fail to correspond to genuine concepts.⁴⁸

The second general fact about concepts is that not every genuine concept has reliable constitutive rules. The most striking examples of this are pejorative concepts and other “thick” normative concepts that involve false presuppositions. Consider, for instance, Dummett’s example of the xenophobic concept *boche*.⁴⁹ “Boche” was a derogatory term used during World War I by French soldiers to refer to Germans. (I don’t use a contemporary racist term for the obvious reasons.) Plausibly, the constitutive rules for *boche* are something like the following:⁵⁰

From *so-and-so is German*, infer *so-and-so is a boche*.
From *so-and-so is a boche*, infer *so-and-so is brutish and uncivilized*.

Racist and xenophobic concepts like *boche* plausibly count as genuine concepts; surely, thinkers have had thoughts involving them. But no matter which semantic value is assigned to *boche*, its constitutive rules do not preserve truth. By employing these rules, thinkers can infer that arbitrary Germans are brutish and uncivilized. The rules are thus grossly unreliable.

There are many other plausible examples of genuine concepts with unreliable constitutive rules. For instance, there are mathematical concepts with inconsistent constitutive rules, such as the naïve concept of set, the naïve concept of area (as applied to arbitrary subsets of the plane), and the naïve concept of probability (which involves

commitments to countable additivity as well as to the claim that there are uniform probability distributions on denumerable sets). There are examples of physical concepts, such as the intuitive concept of temperature and the Newtonian concept of mass. There are examples of semantic concepts, such as Frege's concept of an extension. Finally, there may even be examples from logic. For example, some intuitionists claim that the classical concept of negation has incoherent constitutive rules. And if truth is a logical concept, the liar paradox demonstrates that the naïve concept of truth is a logical concept with inconsistent constitutive rules.

The second general fact about concepts already shows that there is something wrong with the proposed answer to the reliability challenge: The semantic principle it relies upon is false. The semantic values of genuine concepts are not always assigned in such a way as to make their constitutive rules reliable. To be extensionally adequate, the semantic principle must somehow be restricted. And this restriction had better not be *ad hoc* and unmotivated. It is not clear what plausibly can be said here.⁵¹

There is a still more serious issue. Each of the two general facts entails that not every conceptual role corresponds to a genuine concept with reliable constitutive rules. Indeed, most conceptual roles do not.

This should not be a surprise; for the case of logic, it follows from our assumption that logic is objective. In particular, it follows from the no-plenitude claim about logic. Only a small number of logical practices are correct, and so very many conceptual roles for logical constants fail to correspond to genuine concepts with reliable constitutive rules.

This result has a significant ramification. Namely, the proposed solution to the reliability challenge provides (at best) *a very partial explanation of our reliability*. Suppose that it turns out that certain conceptual roles – call them the “good roles” – are guaranteed by the principles of semantics to correspond to genuine concepts with reliable constitutive rules. The concept-constitution explanation can explain how it is that our deductive rules are reliable given that they help to constitute good conceptual roles. But an explanation is still owed of how it is that we have good conceptual roles. This is a significant challenge in its own right. To accommodate the two general facts about concepts presented above, as well as the no-plenitude thesis, the restriction to good roles must be a very significant restriction. So the proposed explanation of our reliability cannot provide the entire story.

The moral, then, is this: There are significant difficulties in getting a concept-constitution account of reliability to work. Moreover, even if these difficulties can be resolved, the reliability challenge for logic will not fully be answered.

Evolution

Where does this leave us? The only apparent remaining strategy is to appeal to some kind of evolutionary account. In what follows, I’ll focus on evolution by natural selection. But my discussion will apply more generally. In particular, it will apply to other mechanisms that involve selection including, for example, certain forms of cultural evolution.

On the most straightforward evolutionary explanation of how we came to employ reliable deductive rules, our ancestors were selected for employing reliable deductive

rules – a heritable trait. This explains how it is that we, their descendants, came to employ reliable deductive rules.⁵²

I believe that an explanation of this general kind will ultimately turn out to be correct. At the very least, it is the most promising avenue available. But it is important to note that such an approach faces significant difficulties. There are several reasons to think that evolutionary accounts cannot even in principle explain how we came to employ reliable deductive rules.

Here are some of the most pressing concerns: (i) Even if evolutionary accounts can explain why it is that we have useful cognitive mechanisms, they cannot explain why it is that we have highly reliable cognitive mechanisms; (ii) Since engaging in deductive reasoning does not yield new information about the world, there could not have been any evolutionary advantage in so doing; (iii) Even if evolutionary accounts can explain how it is that we employ deductive rules that are actually truth-preserving concerning a narrow range of simple propositions – food, danger, shelter, reproduction, and so on – they seem unable to explain the full extent of our reliability. For example, our deductive rules preserve truth when applied to propositions of arbitrary complexity. They preserve truth when applied to propositions with arbitrary subject matters. And they necessarily preserve truth. It is difficult to see how there could be any selection pressure for the employment of rules with these features.

I believe that there are answers to each of these concerns. In answer to the first objection, possessing highly reliable general-purpose reasoning mechanisms is useful for survival-enhancing tasks such as problem solving and planning for future contingencies. In answer to the second objection, the evolutionary advantage of engaging in deductive

reasoning is presumably not that it provides new information about the world, but that it enables thinkers to convert information into a more usable form. In answer to the third objection, the explanation of why we employ deductive rules that are necessarily truth-preserving when applied to propositions of arbitrary complexity and with arbitrary subject matters is presumably that this is a byproduct of the trait for which our ancestors were really selected – employing an efficient cognitive mechanism that is truth-preserving for a limited range of propositions. But spelling out the details requires a very long story.⁵³

In the remainder of this paper, I would like to discuss one other objection, since this will lead to a final refinement of the reliability challenge. The objection is that evolutionary accounts can only explain why a particular trait predominates in a population and not why particular individuals have the traits that they do. It is a striking fact that particular individuals have reliable mechanisms, and so evolutionary accounts seem incapable of explaining everything that needs to be explained.

To illustrate this worry, it is helpful to consider a toy example. Suppose that there is a bag containing many different colored marbles. Someone (perhaps imperfectly) selects the blue marbles, so that most of the blue ones are retained and most of the rest are thrown out. To make the example more fanciful, we can imagine that the marbles periodically reproduce, with each baby marble having the same color as its parent.

If asked why most of the current population of marbles is blue, one can easily provide an explanation involving selection: There were originally marbles of many different colors, the blue ones were selected, and they passed on their color to their descendants. In contrast, the explanation of why some particular marble is blue does not

involve selection: The marble is blue because it had a blue parent and blueness is hereditary.

This example concerns artificial selection and not natural selection. Yet, the point generalizes. Natural selection can explain why organisms with a phenotypic trait came to predominate in a population. It cannot explain why a particular organism has the trait in question.⁵⁴ The explanation of this latter fact depends on other considerations, namely heredity and random mutation.

This feature of selection has an important consequence for the case of interest here – where the phenotypic trait is that of employing a reliable deductive rule. Consider some particular thinker – Bob – who employs a reliable deductive rule. Assuming an evolutionary account, the explanation of why Bob employs the rule in question does not involve selection, but rather heredity and mutation. It is irrelevant to this explanation that the rule is reliable. That Bob employs a reliable rule is merely an accident. If it was not for the occurrence of some highly chancy events, Bob would not have employed a reliable rule.

This is what motivates the objection to evolutionary responses to the reliability challenge. It seems plausible that to fully respond to the reliability challenge, we must both (i) provide a satisfying explanation of why the population primarily includes thinkers who employ reliable deductive rules and (ii) provide a satisfying explanation of why particular thinkers employ reliable rules. Even if an evolutionary account can satisfy the first constraint, it cannot satisfy the second. And it would be highly unsatisfying to be forced to claim that it was merely a random accident that each individual thinker employs reliable deductive rules. Or so goes the objection.

While this objection is initially compelling, it is ultimately misguided. The striking fact in need of explanation is the population-level fact: Our population primarily includes thinkers who employ reliable deductive rules. That a particular individual employs reliable deductive rules is not nearly as striking or in need of explanation.

To see this, it may help to think about a different scenario. Suppose that there was a vast population of heterogeneous thinkers, each employing a different set of inferential rules. Suppose that very few of these thinkers employed reliable deductive rules. Suppose, in particular, that the number of such thinkers was roughly what one would expect if the inferential rules were distributed by some kind of random process. If the reliable individuals had no other striking properties in common, we would not think that their reliability was particularly in need of explanation. Nor would we find it troubling if it turned out to be merely an accident that those particular individuals were reliable.

The correct response to the objection, then, is to concede that evolutionary explanations can only explain population-level facts. But that is all they need to explain. All that is needed is for there to be an answer to the following question:

The Reliability Challenge (3): How is it that our population predominately includes thinkers with reliable cognitive mechanisms for deductive inference?

This is the right way to understand the reliability challenge for logic. This is the crux of the reliability challenge, whether or not an evolutionary explanation can ultimately be made to work.

6. Conclusion

Let's take stock. As we have seen, the reliability challenge for a domain can be understood as the challenge of explaining (i) how it is that the relevant cognitive

mechanisms work such that they are reliable; and (ii) how it is that we have reliable cognitive mechanisms. If we had good reason to doubt that there were plausible explanations, this would put pressure on the claim that we are reliable about the domain, the claim that the domain is objective, or on our general background views about the world. For the cases of logic, mathematics, modality, and other *a priori* domains, there is a straightforward answer to the first part of the challenge. Our cognitive mechanisms are reliable because they depend on the employment of necessarily truth-preserving rules. However, as we have seen, there is no straightforward answer to the second part of the challenge. Indeed, two seemingly attractive responses – appealing to rational insight and appealing to the nature of concept possession – fail to meet the challenge.

For the case of logic, there is a promising avenue to pursue – namely, an evolutionary approach to explaining our reliability. This approach faces some significant difficulties, but they are difficulties that I think can be overcome.

What about other *a priori* domains, for instance mathematics, modality, and morality? Here, the prospects of an evolutionary approach are less promising.

For the case of mathematics, there may be an evolutionary explanation of how we came to possess particular mathematical concepts – for instance, the concepts *shape* and *number*. There may also be an evolutionary explanation of how we came to have consistent mathematical theories; there was presumably some evolutionary pressure to avoid inconsistency. (Indeed, there was presumably some pressure to avoid theories that are not conservative extensions of our non-mathematical theories and practices.) But assuming that mathematics is objective, consistency (or conservativeness) does not

suffice for truth.⁵⁵ So it is extremely difficult to see how evolution could explain how we ended up with true mathematical theories.

One way to close this gap would be to adopt a plenitudinous view of mathematics.⁵⁶ For instance, one could claim that mathematical reality is so large and multifarious that any consistent mathematical theory truly describes some portion of it. Adopting this view has a serious cost; it requires giving up on the objectivity of mathematics. Indeed, while a plenitudinous view is intuitively plausible for certain branches of mathematics – for instance, abstract algebra – it is not very plausible for others – for instance, arithmetic and set theory. We have strong intuitions of objectivity concerning the natural numbers and (to a lesser extent) the sets. Arithmetic seems special; we don't think that it is just another arbitrary mathematical theory on a par with very many others.

For the case of alethic modality, a broadly evolutionary story may again have a partial role to play. On a now popular view, our reliability in reasoning about necessity and possibility is a side effect of our reliability in reasoning with counterfactuals.⁵⁷ Supposing that this is correct, evolution may help to explain the reliability of some of our modal reasoning. It may help to explain how we reason correctly about counterfactuals concerning certain ways the world might have been – namely, the nearby, “live” possibilities. It is clearly advantageous, for instance, for a thinker to be able to reason about what he could do if a predator were approaching. Yet, additional theoretical resources are needed to explain the full extent of our reliability. One worry here is that in attributing necessity to a proposition, a thinker claims that the proposition would have been true even if the world were radically different. It is difficult to see how there could

have been any evolutionary pressure to reason reliably about radically different possible worlds.⁵⁸ A second, related, worry concerns possibility: It is difficult to see how there could have been any evolutionary pressure to correctly identify which of the many descriptions of outlandish scenarios describe genuine metaphysical possibilities.

In response to these concerns, there are two initially promising options that may be pursued. The first option is to adopt a non-objective view of necessity and possibility. For example, one could adopt a plenitudinous view of modality or a view on which our practices somehow constitute the modal facts. The second option is to claim that our reliability about attributions of necessity and possibility is a byproduct of some other cognitive abilities that were evolutionarily advantageous, such as our ability to reason about live possibilities. Yet, each of these options is problematic. It is difficult to state a non-objective view of necessity and possibility without relying on a modal operator – for instance, in specifying which modal theories correctly describe portions of the plenitudinous modal reality or in specifying which practices can generate modal facts. It is also difficult to see how our reliability about radically different possible worlds could be a byproduct of our reliability about live possibilities.

Finally, consider the case of morality. This is, perhaps, the most troublesome case. It is difficult to see how an evolutionary story can answer the reliability challenge for morality. The trouble is that the connection between evolutionary fitness and correctly reasoning about morality seems very tenuous. Assuming an objective view of morality, it is difficult to see why there would be any evolutionary pressure to have true moral beliefs or to reason correctly about moral issues.⁵⁹ A natural suggestion to make is that acting in moral ways is conducive to survival and flourishing.⁶⁰ But it is not clear that this is so; we

are familiar with numerous ways in which acting morally can hinder the pursuit of our projects.

The prospects for an evolutionary answer to the reliability challenges for mathematics, modality, and morality thus seem mixed. This suggests that we should draw three general morals. First, there may not be a unified answer to the reliability challenge across *a priori* domains. The explanation of our reliability about logic, for example, may be very different from the explanation of our reliability about mathematics.⁶¹ Second, the epistemology of logic may be in better shape than the epistemology of other *a priori* domains. And third, the reliability challenge for these domains remains very pressing.⁶²

Notes

¹ Of course, there are logical truths that are too complex to be easily expressed in natural language or believed by ordinary thinkers.

² Propositions here must be understood to be fine-grained. That is, there can be distinct logically equivalent propositions. It would serve my purposes equally well to rely on a framework in which the primary bearers of logical truth and logical falsity are sentences, so long as sentences are (in part) semantically individuated.

³ See Schechter and Enoch (2006) and Enoch and Schechter (2008) for discussion of the challenge of explaining our epistemic responsibility.

⁴ This characterization of our reliability about logic echoes the characterization of our reliability about mathematics found in Field (1989). Thanks to Paul Boghossian for helpful discussion. Additional hedges may be needed for a correct statement of our reliability. If so, the reader should take them to be implicit. The characterization of our reliability might also be strengthened by adding a partial converse: We believe most simple logical truths and disbelieve most simple logical falsehoods (upon reflection and discussion), on some intuitive notion of simplicity.

⁵ This way of stating the reliability challenge for logic is somewhat artificial. A more natural formulation will emerge below.

Hale (1994) presents a reliability challenge for logic. But he does not view the challenge as worrisome. Indeed, he argues that the existence of a reliability challenge for logic shows that the analogous challenge for mathematics is not serious. Field (2005) argues that there is no viable reliability challenge for logic.

⁶ The reliability challenge for a domain can be thought of as a part of the general “Integration Challenge” of reconciling the metaphysics and epistemology of the domain, as presented in Peacocke (1999).

⁷ Those skeptical that there is a well-behaved distinction between the *a priori* and the *a posteriori*, such as Williamson (2008), can understand the general challenge in terms of the class of armchair propositions – those propositions that can be known from the armchair.

⁸ See the introduction and title essay of Field (1989). See Benacerraf (1973) for an influential precursor.

⁹ See Mackie (1977), page 38. Mackie’s argument is often presented as though it relies purely on metaphysical considerations. But he emphasizes that an important strand of his argument is epistemic. Mackie’s discussion is stated in terms of the difficulties of explaining moral knowledge. So it is unclear whether he thinks the issue primarily concerns the descriptive property of reliability or a normative property such as justification.

Enoch (2010) presents a version of the reliability challenge for the normative realm. Street (2006) can be understood as presenting a reliability challenge for morality.

¹⁰ This is perhaps part of what Kreisel meant by his famous dictum that the central problem in the philosophy of mathematics is not the existence of mathematical objects but the objectivity of mathematical discourse. Kreisel’s discussion of objectivity is scattered throughout his writings. For an example, see Kreisel (1958), page 138.

¹¹ See Wright (1992).

¹² This is a variant of the first component of “minimal realism” in Rosen (1994), page 280.

¹³ To use the terminology of Fine (2001), our thoughts, language, and social practices do not partly ground the logical facts. See Jenkins (2008), chapter one, for the proposal that realism should in general be understood in constitutive terms.

¹⁴ The no-plenitude thesis rules out the claim that every non-explosive logical practice is correct. (A logic is explosive if every proposition logically follows from any proposition.) The no-plenitude thesis also rules out the logical fictionalism of Akiba (2000), according to which any logical practice is correct so long as it is conservative over the atomic facts and the non-logical entailment relations among atomic facts.

The no-plenitude thesis does not obviously rule out the logical pluralism of Beall and Restall (2006). On their view, there is a single set of logical connectives and a single set of true object-level sentences (that is, sentences that do not contain metalogical vocabulary). Where they are pluralists is at the meta-level: They claim that there are many different logical consequence relations. In effect, Beall and Restall are pluralists about metalogic – attributions of logical truth and logical entailment – but not about logic proper.

¹⁵ This view is known as “Plenitudinous Platonism” or “full-blown Platonism”. See Balaguer (1998). See Linsky and Zalta (1995) for a related view.

¹⁶ Field (1998) similarly claims that Plenitudinous Platonism should count as a non-objective view.

¹⁷ There is a fifth option: The practices in use are so plenitudinous that some community was bound to get it right. This is not a promising avenue to explore for two reasons. First, it is false that there is a vast plenitude of logical practices in use. Second, even if there were many such practices, the proposed explanation at best accounts for the fact that *someone* is reliable. It does not account for the stronger fact that *we* are reliable.

¹⁸ The relevant background views concern the likelihood of certain outlandish scenarios, the *a priority* of the domain, and the nature of explanation. In particular, they include our views about when explanation is necessary and what counts as a satisfying explanation.

¹⁹ There may be logical truths that are not believed on the basis of reasoning. One candidate is the Law of the Excluded Middle, since the standard Natural Deduction proof of it is less intuitive than the principle itself.

²⁰ This is a common view. See, for example, Field (2000), Goldman (1986), Pollock and Cruz (1999), and Wedgwood (2002). It shouldn't be pretended that there are no obscurities in a rule-based picture of reasoning. See Kripke (1982) and Boghossian (1989; 2008) for some of the difficulties. But I am not aware of any alternative picture that can do the same work in explaining the nature of reasoning.

²¹ See Gentzen (1934/1935), page 74.

²² Harman (1988; 1995), among others, has made a convincing case that we do not and ought not to employ deductive rules as simple as the ones listed. Harman makes the further claim that we do not reason with deductive rules at all. He does not provide support for this stronger claim. Moreover, it is intuitively implausible. When we report or rehearse our trains of thought, it seems evident that we make use of deductive rules.

²³ There is disagreement in the psychological literature about the correct view of deductive reasoning. The view assumed here is closest to that of Rips (1994) and Braine and O'Brien (1998). Johnson-Laird and Byrne (1991) argue that deductive reasoning is based on reasoning with "mental models". My discussion would only have to be minimally changed if this view were adopted. Other psychologists claim that we do not employ topic-neutral rules of inference, but only domain-specific reasoning mechanisms. See Cheng and Holyoak (1985) and Cosmides (1989) for prominent examples. There are reasons to be leery of such a view. Moreover, adopting such a view would not lessen the challenge of explaining our reliability about logic. We would have to explain the reliability of each of the relevant mechanisms. And there would be the additional challenge of explaining how we arrive at general logical truths using only domain specific mechanisms.

²⁴ See Williamson (2008) for an analogous view on which our ability to reason correctly about metaphysical necessity and possibility is a byproduct of our more basic ability to reason with counterfactuals. See Hill (2006) for a similar proposal.

²⁵ There is also the related challenge of explaining how it is that we have a cognitive mechanism for deductive reasoning with sufficient inferential power to enable us to infer a wide range of logical truths.

²⁶ See Lewis (1986), pages 108–115. It is somewhat difficult to interpret this passage. My reconstruction is close to that of Field (1989), page 233. Lewis's primary interest, of course, concerns the epistemology of extreme modal realism.

²⁷ See Katz (1995), pages 505–506, for a related view. The line of thought is also related to what Linnebo (2006) has called the “Boring Explanation”.

²⁸ Field (1989), pages 233–239, presents several objections to Lewis’s argument.

²⁹ An answer to an etiological question seems to be what Gibbard (2003), pages 253–258, calls a “deep vindication”.

As stated, the etiological question is ambiguous. It can be understood as asking a question about ontogeny – how it is that the relevant creatures develop a reliable mechanism – or about phylogeny – how it is that the mechanism (or the developmental plan for it) is prevalent in the relevant population. The latter is the philosophically more pressing question.

The etiological question may be slightly misnamed. An answer to an etiological question need not provide the causal history of the relevant cognitive mechanism. For instance, one potential answer to the etiological question about logic is that cognitive mechanisms were initially randomly distributed in the population and then some deity decided to smite those individuals lacking reliable deductive mechanisms. (I owe this example to Jonathan Ichikawa.) Indeed, in the case of a non-objective domain, an answer to the etiological question need not provide a causal history at all. For instance, a different potential answer to the etiological question is that the facts of the domain are constituted by our opinions, and so we would have possessed a reliable cognitive mechanism no matter what. In the case of an objective domain, however, it is difficult to envision an answer that does not involve some kind of causal story.

³⁰ This is not to say that there are no interesting empirical questions concerning the operation of our cognitive mechanism for deductive inference.

³¹ Lewis is mistaken in his claim that the relevant distinction is between necessary and contingent domains. There is no quick answer to the operational question for the familiar necessary *a posteriori* truths. More interestingly, there is a quick answer for certain contingent truths. For instance, there is a straightforward explanation of how I am reliable in my belief that if *p* obtains then *p* actually obtains. In general, there is a straightforward answer to the operational question in cases where the relevant cognitive method is guaranteed to output a truth given true inputs. This may happen even if the content of the output varies from world to world due to the presence of an indexical. Thanks to Tim Williamson for helpful discussion of this issue.

³² See Field (1989), pages 26–27.

³³ It is difficult to provide a general account of strikingness. See Horwich (1982) and Schlesinger (1991) for attempts. I take it that the most promising proposals are: (i) A phenomenon is striking if it can be described using a simple rule; and (ii) A phenomenon is striking if there is a salient theory that would predict or explain it.

³⁴ Thanks to Ralph Wedgwood for pressing this line of thought.

³⁵ Here’s a second kind of example: Any contingent fact will logically entail every logical truth. But the explanation of a contingent fact will not be an explanation of arbitrary logical truths.

³⁶ This is a modified version of the argument for immodest inductive rules in Lewis (1971). Also see Field (2000) and Elga (2010).

³⁷ For versions of the second and third grounds, see Nagel (1997), chapter four.

³⁸ There are, however, reasons for doubt. For instance, we should in principle be open to arguments for alternative logics.

³⁹ As Rowland Stout pointed out to me, there is a fourth proposal to consider – namely, learning. One way to develop this suggestion is to claim that we acquired our deductive competence via trial and error learning. We test logically complex propositions against the world (via perception) and use such tests to evaluate the truth-conduciveness of proposed deductive rules. The problem with this view is that, while it may provide the correct account of some of our deductive reasoning, it is implausible for the core parts of deduction. Moreover, even if the proposal were true, it would simply raise another explanatory demand. We would have to answer the reliability challenge for our cognitive mechanism for trial and error learning. This mechanism could not itself be acquired by trial and error learning. In general, if the reliability challenge for a domain is answered by appeal to more basic cognitive mechanisms, we can simply raise reliability challenges for those mechanisms.

⁴⁰ Gödel (1947) may have endorsed a rational insight-based view of mathematics.

BonJour (1998) endorses a rational insight-based view of *a priori* knowledge in general.

⁴¹ See, for instance, Boghossian (2003), pages 230–232. Also see Wright (2004) for objections particular to the case of logic.

⁴² See Bealer (1996) for the claim that we have intellectual seemings. See Sosa (1996) for arguments that these mental states are best understood as dispositions to believe.

⁴³ My development of this proposal makes use of the work of Bealer (1996; 1999) and Peacocke (1992; 1993), although the view described here should not be ascribed to either of them.

⁴⁴ Bealer claims that what is constitutive of possessing concepts is that we have certain intuitions, not that we employ certain rules. This difference does not matter for my purposes. Bealer also suggests that what is constitutive of possessing concepts is that we have certain reliable intuitions. If that is his view, appealing to concept constitution does no work in answering the reliability challenge. An explanation is needed of why we have reliable intuitions.

⁴⁵ See Block (1997) for a discussion of Conceptual Role Semantics. See Fodor and Lepore (1991) and Williamson (2003) for objections.

⁴⁶ See Prior (1960).

⁴⁷ Other examples of this kind include “vel” as described in Peacocke (2004), page 18, and “plonk” as described in Belnap (1962).

⁴⁸ This is one place that a notion of harmony – the idea that there should be some sort of match between introduction and elimination rules – might be deployed. See Tennant (1987) and Dummett (1991). However, it has proved difficult to provide a plausible general account of harmony that extends beyond the case of logical concepts.

⁴⁹ The example of *boche* is originally due to Dummett (1973), pages 397 and 454.

Boghossian (2003) makes use of this example to argue that concept-constituting rules need not be truth-preserving.

⁵⁰ See Horwich (2005), pages 153–154, Williamson (2003), pages 257–268, and Williamson (2009) for objections to the claim that these are the constitutive rules of a genuine concept.

- ⁵¹ This is a second place where a notion of harmony might be deployed.
- ⁵² This evolutionary account could be complicated in many ways. For instance, it could be combined with a story on which trial and error learning is important. Or it could be combined with some version of the concept-constitution view discussed above. For my purposes here, it suffices to consider only the simplest possible story.
- ⁵³ See Schechter (MS) for an attempt to address these difficulties.
- ⁵⁴ See Sober (1984), section 5.2. As Neander (1995) points out in response, repeated bouts of selection can help to explain why it is likely that a trait would emerge in a population. However, this does not affect the point in the text, since selection cannot explain why a particular individual has the trait in question.
- ⁵⁵ See Field (1989).
- ⁵⁶ See Balaguer (1998) and Linsky and Zalta (1995).
- ⁵⁷ See Williamson (2008) and Hill (2006).
- ⁵⁸ See Nozick (2001), page 122, and Stroud (2000) for versions of this worry.
- ⁵⁹ Street (2006) has argued on related grounds for a non-objective constructivist view of morality.
- ⁶⁰ See Enoch (2010) for an explanation of our reliability about the normative with this general shape.
- ⁶¹ There may not even be a uniform explanation of our reliability about logic. As Alex Paseau reminded me, parts of second-order logic closely resemble mathematics.
- ⁶² Versions of this material were presented at an NYU dissertation seminar, the Arché core seminar at the University of St. Andrews, a philosophy of mathematics seminar and a discussion group at Oxford University, and a colloquium at Bristol University. I would like to thank the audiences at these events for their helpful questions and comments. I would also like to thank Paul Boghossian, Jessica Brown, Ray Buchanan, Winston Chiong, David Christensen, David Enoch, Greg Epstein, Dana Evan, Hartry Field, Kit Fine, Don Garrett, Pete Graham, Liz Harman, Jonathan Ichikawa, Dan Isaacson, Øystein Linnebo, Anna-Sara Malmgren, Alex Paseau, Christopher Peacocke, Jim Pryor, Erica Roedder, Karl Schafer, Jonathan Schaffer, Brad Skow, Declan Smithies, Sharon Street, Rowland Stout, Michael Strevens, Scott Sturgeon, Ralph Wedgwood, Roger White, Tim Williamson, Crispin Wright, and Masahiro Yamada for helpful discussion.

References

- Akiba, Ken. (2000) "Logic and Truth: A Fictionalist View," *Journal of Philosophical Research*, 25, 101–123.
- Balaguer, Mark. (1998) *Platonism and Anti-Platonism in Mathematics*. Oxford: Oxford University Press.
- Bealer, George. (1996) "A Priori Knowledge and the Scope of Philosophy," *Philosophical Studies*, 81, 121–142.
- Bealer, George. (1999) "A Theory of the A Priori," *Philosophical Perspectives*, 13, 29–55.
- Beall, J.C. and Greg Restall. (2006) *Logical Pluralism*. Oxford: Oxford University Press.
- Belnap, Nuel. (1962) "Tonk, Plonk and Plink," *Analysis*, 22, 130–134.

- Benacerraf, Paul. (1973) "Mathematical Truth," *The Journal of Philosophy*, 70, 661–679.
- Block, Ned. (1997) "Semantics, Conceptual Role," *Routledge Encyclopedia of Philosophy*.
- Boghossian, Paul. (1989) "The Rule-Following Considerations," *Mind*, 98, 507–549.
- Boghossian, Paul. (2003) "Blind Reasoning," *Proceedings of the Aristotelian Society, Supplementary Volume*, 77, 225–248.
- Boghossian, Paul. (2008) "Epistemic Rules," *The Journal of Philosophy*, 105, 1–29.
- BonJour, Laurence. (1998) *In Defense of Pure Reason*. Cambridge: Cambridge University Press.
- Braine, Martin, and David P. O'Brien. (eds.) (1998) *Mental Logic*. Mahweh, NJ: Erlbaum.
- Cheng, Patricia and Keith Holyoak. (1985) "Pragmatic Reasoning Schemas," *Psychology*, 17, 391–416.
- Cosmides, Leda. (1989) "The Logic of Social Exchange," *Cognition*, 31, 187–276.
- Dummett, Michael. (1973) *Frege: Philosophy of Language*. London: Duckworth.
- Dummett, Michael. (1991) *The Logical Basis of Metaphysics*. Cambridge, MA: Harvard University Press.
- Elga, Adam. (2010) "How to Disagree about How to Disagree," in Richard Feldman and Ted A. Warfield (eds.), *Disagreement* (pp. 175–186). Oxford: Oxford University Press.
- Enoch, David. (2010) "The Epistemological Challenge to Metanormative Realism," *Philosophical Studies*, 48, 413–438.
- Enoch, David and Joshua Schechter. (2008) "How Are Basic Belief-Forming Methods Justified?," *Philosophy and Phenomenological Research*, 76, 547–579.
- Field, Hartry. (1989) *Realism, Mathematics and Modality*. Oxford: Basil Blackwell.
- Field, Hartry. (1998) "Which Undecidable Mathematical Sentences Have Determinate Truth Values?," in Harold Dales and Gianluigi Oliveri (eds.), *Truth in Mathematics* (pp. 291–310). Oxford: Oxford University Press.
- Field, Hartry. (2000) "Apriority as an Evaluative Notion," in Paul Boghossian and Christopher Peacocke (eds.), *New Essays on the A Priori* (pp. 117–149). Oxford: Oxford University Press.
- Field, Hartry. (2005) "Recent Debates about the A Priori," *Oxford Studies in Epistemology* 1, 69–88.
- Fine, Kit. (2001) "The Question of Realism," *Philosophers' Imprint*, 1, 1–30.
- Fodor, Jerry, and Ernie Lepore. (1991) "Why Meaning (Probably) isn't Conceptual Role," *Mind and Language*, 6, 328–34.
- Gentzen, Gerhard. (1934–35) "Untersuchungen über das Logische Schliessen," *Mathematische Zeitschrift*, 39, 176–210 and 405–431.
- Gibbard, Allan. (2003) *Thinking How to Live*. Cambridge, MA: Harvard University Press.
- Gödel, Kurt. (1947) "What is Cantor's Continuum Problem?," in Solomon Feferman, et al. (eds.), *Collected Works of Kurt Gödel*, Volume II (pp. 176–188) Oxford: Oxford University Press.
- Goldman, Alvin. (1986) *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.

- Hale, Bob. (1994) "Is Platonism Epistemologically Bankrupt?," *The Philosophical Review*, 103, 299–325.
- Harman, Gilbert. (1988) *Change in View: Principles of Reasoning*. Cambridge, MA: The MIT Press.
- Harman, Gilbert. (1995) "Rationality," in Edward Smith and Daniel Osherson (eds.), *Thinking: An Invitation to Cognitive Science*, volume 3, second edition (pp. 175–211). Cambridge, MA: The MIT Press.
- Hill, Christopher. (2006) "Modality, Modal Epistemology, and the Metaphysics of Consciousness," in Shaun Nichols (ed.), *The Architecture of the Imagination* (pp. 205–235). Oxford: Oxford University Press.
- Horwich, Paul. (1982) *Probability and Evidence*. Cambridge: Cambridge University Press.
- Horwich, Paul. (2005) "Meaning Constitution and Epistemic Rationality," chapter 6 in his *Reflections on Meaning*. Oxford: Oxford University Press.
- Jenkins, Carrie. (2008) *Grounding Concepts: An Empirical Basis for Arithmetic Knowledge*. Oxford: Oxford University Press.
- Johnson-Laird, Philip, and Ruth Byrne. (1991) *Deduction*. Sussex: Lawrence Erlbaum Associates.
- Katz, Jerrold. (1995) "What Mathematical Knowledge Could Be," *Mind*, 104, 491–522.
- Kreisel, Georg. (1958) "Wittgenstein's Remarks on the Foundations of Mathematics," *The British Journal for the Philosophy of Science*, 9, 135–158.
- Kripke, Saul. (1982) *Wittgenstein on Rules and Private Language*. Cambridge, MA: Harvard University Press.
- Lewis, David. (1971) "Immodest Inductive Methods," *Philosophy of Science*, 38, 54–63.
- Lewis, David. (1986) *On the Plurality of Worlds*. Oxford: Blackwell Publishers.
- Linnebo, Øystein. (2006) "Epistemological Challenges to Mathematical Platonism," *Philosophical Studies*, 129: 545–574.
- Linsky, Bernard and Edward Zalta. (1995) "Naturalized Platonism versus Platonized Naturalism," *The Journal of Philosophy*, 92, 525–555.
- Mackie, John L. (1977) *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin.
- Nagel, Thomas. (1997) *The Last Word*. Oxford: Oxford University Press.
- Neander, Karen. "Pruning the Tree of Life," *The British Journal for the Philosophy of Science*, 46, 59–80.
- Nozick, Robert. (2001) *Invariances: The Structure of the Objective World*. Cambridge, MA: Harvard University Press.
- Peacocke, Christopher. (1992) *A Study of Concepts*. Cambridge, MA: The MIT Press.
- Peacocke, Christopher. (1993) "How are A Priori Truths Possible?," *European Journal of Philosophy*, 1, 175–199.
- Peacocke, Christopher. (1999) *Being Known*. Oxford: Oxford University Press.
- Peacocke, Christopher. (2004) *The Realm of Reason*. Oxford: Clarendon Press.
- Pollock, John and Joseph Cruz. (1999) *Contemporary Theories of Knowledge*. Lanham: Rowman and Littlefield.
- Prior, Arthur. (1960) "The Runabout Inference Ticket," *Analysis*, 21, 38–39.
- Rips, Lance. (1994) *The Psychology of Proof: Deductive Reasoning in Human Thinking*. Cambridge, MA: The MIT Press.

- Rosen, Gideon. (1994) "Objectivity and Modern Idealism: What is the Question?," in Michaelis Michael and John O'Leary Hawthorne (eds.), *Philosophy in Mind: The Place of Philosophy in the Study of Mind* (pp. 277–230). Dordrecht: Kluwer Academic Publishers.
- Schechter, Joshua. (MS) "Could Evolution Explain Our Reliability about Logic?"
- Schechter, Joshua and David Enoch. (2006) "Meaning and Justification: The Case of Modus Ponens," *Noûs*, 40, 687–715.
- Schlesinger, George. (1991) *The Sweep of Probability*. South Bend: University of Notre Dame Press.
- Sober, Elliott. (1984) *The Nature of Selection*. Cambridge, MA: The MIT Press.
- Sosa, Ernest. (1996). "Rational Intuition: Bealer on its Nature and Epistemic Status," *Philosophical Studies*, 81, 151–162.
- Street, Sharon. (2006) "A Darwinian Dilemma for Realist Theories of Value," *Philosophical Studies*, 127, 109–166.
- Stroud, Barry. (2000) "Evolution and the Necessities of Thought," in his *Meaning, Understanding, and Practice* (pp. 52–67). Oxford: Oxford University Press.
- Tennant, Neil. (1987) *Anti-Realism and Logic*. Oxford: Clarendon Press.
- Wedgwood, Ralph. (2002) "Internalism Explained," *Philosophy and Phenomenological Research*, 65, 349–369.
- Williamson, Timothy. (2003) "Understanding and Inference," *Proceedings of the Aristotelian Society, Supplementary Volume*, 77, 249–293.
- Williamson, Timothy. (2008) *The Philosophy of Philosophy*. Oxford: Wiley-Blackwell.
- Williamson, Timothy. (2009) "Reason, Inference and the Semantics of Pejoratives," in Joseph Almog and Paolo Leonardi (eds.), *The Philosophy of David Kaplan* (pp. 137–158). Oxford: Oxford University Press.
- Wright, Crispin. (1992) *Truth and Objectivity*. Cambridge, MA: Harvard University Press.
- Wright, Crispin. (2004) "Intuition, Entitlement and the Epistemology of Logical Laws," *Dialectica*, 58, 155–175.