Could Evolution Explain Our Reliability about Logic?*

Joshua Schechter

1. INTRODUCTION

Let the "logical propositions" be the logical truths and logical falsehoods. We are reliable about logic in the following sense: The logical propositions that we believe (upon reflection and discussion) are by-and-large true and the logical propositions that we disbelieve (upon reflection and discussion) are by-and-large false.[1] This is a striking fact about us, one that stands in need of explanation. But it is not at all clear how to explain it. So we have a puzzle: How is it that our logical beliefs match the logical facts? How is it that we are reliable about logic?

This puzzle is akin to the well-known Benacerraf-Field problem for mathematical Platonism.[2] According to that argument, mathematical Platonists are unable to explain our reliability about mathematics due to their claim that mathematical objects are abstract. In the absence of some amazing cosmic accident, it is difficult to see how we could have ended up with the correct mathematical beliefs and practices. This provides reason to reject Platonism about mathematics.

Logic does not have – or at least does not obviously have – a distinctive ontology. Nevertheless, the challenge of explaining our reliability about logic is also daunting. What gives this reliability challenge its bite is not the ontology of logic but the (apparent)

---

* This is the penultimate draft of a paper forthcoming in Tamar Szabo Gendler and John Hawthorne (eds.) *Oxford Studies in Epistemology* 4.

[1] This characterization of our reliability about logic is broadly analogous to the characterization of our reliability about mathematics in Field (1989).

[2] See the introduction and title essay of Field (1989). See Benacerraf (1973) for an important precursor.

objectivity of logic.[3] We can understand the claim that logic is objective as the conjunction of the following three theses:[4] First, certain sentences and mental representations express logical truths and logical falsehoods. They are therefore both meaningful and truth-apt.[5] Second, 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. Third, even if there is, in some sense, a plenitude of incompatible logical practices, only one – or a small number of them – is correct.[6]

Given the objectivity of logic, it is not at all clear how to explain our reliability. The intuitive difficulty is this: We have some understanding of how we could have veridical beliefs about non-objective facts. For instance, it is not at all mysterious how Sir Arthur Conan Doyle could be reliable about what is true in the Sherlock Holmes fiction. We also have some understanding of how we could arrive at veridical beliefs about objective facts via perception. But this understanding does not extend to the case of logic. We do not understand how we could be reliable about objective facts that were not learned via some kind of perception.

This difficulty is potentially very significant. If we were to come to believe that there is no satisfying explanation of our reliability compatible with the objectivity of logic, this would put pressure on our belief that logic is objective, on our belief that we

[3] The same holds true for mathematics. What gives the Benacerraf-Field problem its bite is not the ontology but the objectivity of mathematics. Indeed, there is a pressing challenge for any domain such that we think (i) we are reliable; (ii) the domain is objective; and (ii) our beliefs about the domain are not generated by some kind of perception.

[4] There may be several different philosophically-interesting notions of objectivity. See Schechter (2010) for defense of the claim that this is the relevant characterization of objectivity in the context of reliability challenges.

[5] To raise a reliability challenge, it suffices that the relevant truth predicate be minimal or deflationary. It is not necessary to appeal to a more robust conception of truth.

[6] Notice that the claim that logic is objective does not entail that it is an objective matter whether a truth counts as a logical truth. That is, logicality need not be objective.

are reliable about logic, or on our general background views about the world. But giving up on the objectivity of logic – or on our reliability or our general background views – would be devastating to our ordinary ways of thinking. Thus, there is an important explanatory challenge to answer.

In this paper, I examine one candidate answer to this challenge. In particular, I pursue the attractive thought that our reliability about logic is to be explained by appeal to evolution by natural selection.[7] The account is based on two main ideas. The first is that our reliability about logic is to be explained by appeal to a more basic competence: We are reliable about logic because we are reliable in our deductive reasoning.[8] The second is that being a reliable deductive reasoner conferred a heritable survival or reproductive advantage upon our ancestors. We inherited this trait. This explains how we are reliable in our deductive reasoning.

In the past, when I have presented this view, it has prompted two radically different responses. Some have claimed that it is obviously correct: We are the products of evolution, and since we are reliable about logic, there must be an evolutionary explanation of this fact. Others have claimed that the view is obviously hopeless. As will become clear, both reactions are premature. There are several in principle difficulties facing evolutionary accounts, but there is reason to believe that they can be addressed.

The difficulties facing evolutionary accounts are fierce: (i) Even if evolutionary accounts can explain why we employ useful cognitive mechanisms, they seem unable to

[7] The idea of explaining the reliability of our reasoning by appeal to evolution has a long history. See, for instance, Darwin (1871), book I, chapter V. Versions of this view can also be found in Spencer, Mach, Avenarius, Boltzmann, Simmel, James, and Dewey, among many others.

[8] Some philosophers and psychologists have claimed that we do not reason deductively but rather employ only inductive or abductive patterns of reasoning. I find this view implausible, but cannot argue against it here. If this alternative view were instead adopted, many of the same issues would arise.

explain why we employ highly reliable ones; (ii) Even if evolutionary accounts can explain why we reason reliably about a narrow range of simple propositions – those concerning danger, food, reproduction, shelter, and so on – they seem unable to explain our reliability concerning propositions of arbitrary complexity and with arbitrary subject matters; (iii) Even if evolutionary accounts can explain how we came to employ deductive rules of inference that are actually truth-preserving, they seem unable to explain how we came to employ rules that are necessarily truth-preserving; Finally, (iv) since engaging in deductive reasoning does not yield novel information about the world, it is difficult to see how there could be any selective advantage in doing so.[9]

The purpose of this paper is to sketch an evolutionary explanation of our reliability about logic and to demonstrate how these general difficulties may be addressed. I do not claim that every detail of my evolutionary explanation is correct. I do not even claim that some evolutionary explanation must be correct. Rather, my main claim is that *there is no in principle reason to think that evolutionary accounts are incapable of explaining why we are reliable about logic*. In particular, the account I sketch provides one plausible answer to the reliability challenge for logic. It demonstrates that there are plausible answers to be had. This defangs the reliability challenge for logic. It defuses the tension between the claim that logic is objective, the claim that we are reliable about logic, and our general background views about the world.

[9] Other alleged difficulties include the following: There is no way to make sense of the distinction between selection for accepting the correct logic whichever that is, and selection for accepting a particular logic that happens to be correct. See Field (1998), page 19. Evolutionary explanations cannot explain why having objective thoughts is within the range of biological options. See Nagel (1986), pages 78–79. The probability that any of our cognitive mechanisms is reliable is low, given naturalism and given that we are the product of evolution, due to the difficulty of providing a naturalistic account of the role of content in causing behavior. See Plantinga (1993), chapter 12.

This paper will proceed as follows. In the next section, I further develop the reliability challenge for logic. I present what I take to be the crux of the challenge – the challenge of explaining how it is that we have a reliable cognitive mechanism for deductive reasoning. In section three, I discuss the nature of evolutionary explanations. Section four is devoted to sketching an evolutionary explanation of our reliability about logic. I first present an explanation of why it is that we possess logical concepts. I then present an explanation of why it is that we employ reliable deductive rules given that we possess these concepts. In so doing, I address the fourth difficulty listed above. Finally, in section five, I answer the remaining three difficulties for evolutionary accounts of our reliability about logic.

2. PRELIMINARIES[10]

Before developing the reliability challenge for logic, I should first explain what I mean by "logic". Logic, as I use the term here, does not concern artificial formal languages. Rather, it concerns propositions that can be expressed in non-technical natural language and believed by ordinary thinkers. Certain propositions are logical truths – for instance, the proposition that every walrus is a walrus. Other propositions are logical falsehoods – for instance, the proposition that some walrus is not a walrus.

There are many interesting questions that arise concerning the nature of logic. For my purposes here, I need not presuppose any particular account. In my discussion, however, I will assume that (i) propositions and not sentences are the primary bearers of logical truth and logical falsity; (ii) propositions are fine-grained in the sense that there

[10] This section presents some of the main conclusions of Schechter (2010).

can be distinct logically equivalent propositions; (iii) whether a proposition is logically true or logically false depends on its logical form; and (iv) logical truths are necessarily true and logical falsehoods are necessarily false, on any reasonable (alethic) kind of necessity. I will also assume that classical logic is at least approximately correct.

As a first pass, the reliability challenge for logic is the challenge of explaining why it is that the logical propositions that we believe (upon reflection and discussion) are by-and-large true and the logical propositions that we disbelieve (upon reflection and discussion) are by-and-large false. Yet, this is not the best way to understand the crux of the challenge. This challenge has a straightforward answer. Our reliability about logic can be explained by appeal to a more basic competence; we are reliable about logic because we are reliable in our deductive reasoning.[11]

Consider some moderately complex logical truth, for instance *if both A and if A then B then B*, substituting particular sentences for A and B. We believe this proposition, at least upon reflection and discussion. How did we come to believe it? There are many possibilities. For instance, we may have learned it from a logic teacher. Or we may have observed that B and inferred the conditional from it. However, the canonical way one comes to believe this truth is via a chain of reasoning, perhaps one like the following:

Suppose both A and if A then B.
    So A.
    So if A then B.
    So B.
So if both A and if A then B, then B.

[11] This claim is broadly analogous to the suggestion that our ability to reason correctly about metaphysical necessity is a byproduct of our more basic ability to reason correctly about counterfactuals. See Hill (2006) and Williamson (2007).

More generally, we come to believe logical truths and disbelieve logical falsehoods on the basis of such deductive reasoning. Our logical beliefs are the outputs of deductive reasoning in cases where there are no initial premises.

It is plausible that whenever thinkers reason in this way, their reasoning depends on the employment of rules of inference. In particular, it is plausible that deductive reasoning relies on the employment of rules that resemble the rules that that appear in natural deduction formulations of logic. On this view, deductive rules include such rules as:

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

There are several reasons to think that the rules that we employ in deductive reasoning are more complex than the simple rules listed above.[12] Nevertheless, it is plausible that the rules that we employ are closely related to the standard natural deduction rules, and in what follows, I'll assume that they are.[13]

Consider again *if both A and if A then B, then B*. We believe this proposition because we went through a chain of reasoning like the one displayed above. We ended up with a true belief because the transitions involved in our reasoning were truth-preserving. The transitions were truth-preserving because the deductive rules that governed them are

[12] See Harman (1988; 1995).
[13] There is disagreement in the psychological literature over the correct view of deductive reasoning. The view assumed here is closest to that of Rips (1994) and Braine and O'Brien (1998). My discussion would have to be changed if an alternative view were adopted. But in most cases, the changes would be minimal.

reliable. In general, we are reliable in our logical beliefs because we are reliable in our deductive reasoning. We employ reliable deductive rules.[14]

Of course, providing this explanation does not fully answer the reliability challenge for logic. It raises a new explanatory demand. Explanation is now needed of how it is that we are reliable in our deductive reasoning. The challenge thus becomes that of explaining the reliability of our cognitive mechanism for deductive reasoning.

What is this challenge? There are two important explanatory questions concerning the reliability of our deductive mechanism. They may be stated as follows:

> *The Operational Question*: How does our cognitive mechanism for deductive reasoning work such that it is reliable?
>
> *The Etiological Question*: How is it that we have a cognitive mechanism for deductive reasoning that is reliable?[15]

To illustrate the distinction between these two questions, it may be helpful to compare a different cognitive mechanism – say, vision. Our visual mechanism is reliable in the sense that it by-and-large produces true beliefs about our environment. This is a striking fact, in need of explanation. Indeed, this fact raises two explanatory questions. First, how does our visual mechanism work such that it reliably produces true beliefs about our environment? Second, how is it that we have a reliable visual mechanism? The answers to these two questions are very different.

To answer the reliability challenge for logic, satisfying answers to both the operational and the etiological questions are needed. The operational question has a

---

[14] There are well-known experimental results showing that humans are prone to errors in deductive reasoning. Perhaps the most striking results concerns variants of the Wason selection test. So the claim that we are reliable in our deductive reasoning should not be overstated. But these errors seem largely to be performance errors. The rules built into our deductive competence are reliable.

[15] The etiological question can be interpreted in two ways – as concerning ontogeny or as concerning phylogeny. To fully answer the reliability challenge, both questions require answers. But the ontogenetic question is philosophically not very pressing. The development of a reliable deductive mechanism is presumably genetically encoded. The philosophically pressing question concerns phylogeny.

straightforward answer: Our deductive mechanism works via the employment of rules of inference. This mechanism is reliable because the deductive rules of inference that we employ are necessarily truth-preserving. They are guaranteed to yield true beliefs from true beliefs.[16] That's all that needs to be said. We need not claim that our deductive mechanism somehow tracks the logical facts.

The etiological question, however, cannot be answered so easily. To answer this question, an explanation is needed of how it is that we have a reliable cognitive mechanism for deductive reasoning. In particular, an explanation is needed of how it is that we employ reliable deductive rules. There does not seem to be a quick explanation of this fact. This is the crux of the reliability challenge for logic.

It is important to recognize that this explanatory demand is not generated by some overly powerful epistemological principle that quickly leads to radical skepticism. For instance, we need not accept the claim that thinkers must possess an explanation of how it is that they have a reliable cognitive method in order to be justified in employing the method. Rather, the motivating line of thought is simply this: It is a striking fact that we have a reliable cognitive mechanism for deductive reasoning. This is a fact that "cries out" for explanation. *Ceteris paribus*, it is a cost of a theory if it treats some striking phenomenon as merely accidental or otherwise inexplicable. Thus, it would be very unsatisfying to be forced to claim that it was a brute fact – an amazing cosmic accident – that we employ reliable deductive rules. If that were the only available account, there would be a tension in our overall view of the world. There would be pressure to somehow

[16] This claim must be generalized to handle deductive rules that involve propositional attitudes other than belief. Let's say that believing a proposition is *correct* if the proposition is true, disbelieving a proposition is correct if the proposition is false, believing a proposition under a supposition is correct if the proposition is true if the supposition is, and so on. Our deductive rules of inference are reliable in the sense that they necessarily preserve correctness. I leave this generalization implicit in what follows.

modify our view. This is why a substantive answer to the etiological question is sorely needed.

3. EVOLUTIONARY EXPLANATIONS

Perhaps the most natural approach to answering the reliability challenge for logic – that is, for answering the etiological question for deductive reasoning – is to appeal to evolution by natural selection. Evolution by natural selection works as follows: There is initially a population of organisms. Each organism has a genetic endowment, its genotype. Each organism also has phenotypic traits that depend in part on its genotype, and so are heritable. There is variation in the genotypes and phenotypic traits among the organisms in the population. Organisms with certain phenotypic traits tend to be better at survival and reproduction than the rest, given their background environment. Over time, these fitter organisms survive longer and reproduce more frequently than less fit organisms. This yields a change in the gene frequency of the population and a corresponding change in the frequency of phenotypic traits.[17]

According to the most straightforward kind of evolutionary explanation of why a phenotypic trait came to predominate in a population, possessing the trait enhanced the fitness of the ancestors of the population (in their background environment). In the case of interest, the relevant trait is that of employing reliable deductive rules. On the most straightforward evolutionary explanation of why our population came to employ reliable deductive rules, then, our ancestors were selected for employing reliable deductive rules

[17] There are biological mechanisms of evolution other than natural selection, such as random drift. I focus on natural selection because such non-selective mechanisms cannot play the central role in a satisfying explanation of our reliability. (Lamarkian inheritance, if it were to exist, could also play such a role.)

– a heritable trait – and this explains why it is that we, their descendants, employ reliable deductive rules.

An evolutionary account of our reliability about deduction is attractive because it is plausible that employing reliable deductive rules can confer survival and reproductive advantages on an organism. For instance, if an organism believes that a predator is either in this grove of trees or that grove, and then comes to learn that the predator is not in that grove, it is advantageous for the organism to come to believe that the predator is in this grove. An evolutionary account is also attractive because it fits well with the general scientific picture of our place in the world.

It may seem strange to discuss evolutionary explanations of traits such as employing *reliable* deductive rules. Whether a rule is reliable is not, in some sense, a purely biological fact. It might seem more promising to consider the trait of employing some particular collection of rules. But this line of thought is mistaken. Employing reliable deductive rules is a phenotypic trait like any other. There is no reason to suppose that evolutionary explanations are unable to explain the presence of such traits. Moreover, if we only provided an explanation of how we came to employ some particular collection of deductive rules, we would not answer the reliability challenge for logic. Even though we would have explained why we employ the particular rules, and even though it is a necessary truth that those rules are reliable, we would not have explained why we employ reliable rules. Explanation is not closed under necessary entailment.[18] The explanation of our employment of the rules would presumably have nothing to do with their reliability. The reliability of our deductive rules would still seem accidental in a worrisome way.

[18] See Schechter (2010), page 447, for an example.

There are several ways in which this simple evolutionary account of our reliability could be modified. For instance, it could be claimed that our ancestors were directly selected not for employing reliable deductive rules, but for possessing linguistic abilities that required or led to the possession of reliable deductive rules. It could be claimed that our ancestors were selected for employing reliable non-deductive learning mechanisms that can be used to acquire reliable deductive rules. It could be claimed that our ancestors were selected to easily adopt reliable deductive rules when taught them by confederates.[19] Or it could be claimed that cultural evolution rather than biological evolution was primary. It is worth emphasizing that it is not part of my task here to reject any of these suggestions. My goal is only to sketch one plausible explanation of the reliability of our deductive rules and show how seemingly powerful difficulties for it may be addressed. The main point of this paper is to provide something like a demonstration of possibility (or of plausibility) and not of actuality.

Before filling in more of the details of the evolutionary account, it is important that it first be made clear what natural selection can – and cannot – explain. To this end, it is helpful to consider a simple toy example.

Suppose that we have very many different sets of fair dice. Let us assume that each set contains five dice and that the sets differ in color. So there are five bright red dice, five faded blue dice, five forest green dice, and so on. Suppose that we roll each of the dice simultaneously and select the sets of dice that come up all sixes. That is, we keep the sets of dice that come up all sixes and dispose of the rest.

[19] This suggestion is related to the Baldwin effect – the idea that there might be selection for organisms that can more easily learn certain abilities and then, only later, for organisms that have those abilities innately.

If asked why the current population contains sets of dice that came up all sixes, we can provide an explanation involving selection: There were very many sets of dice and the ones that came up all sixes were selected. In contrast, there is no selective explanation of why some particular set of dice – say the faded blue dice – came up all sixes. That was purely a matter of chance.

This toy example involves artificial selection, not natural selection. Yet, the point generalizes. Natural selection can help to explain why organisms with a certain phenotypic trait came to predominate in a population. It cannot explain why a particular organism has the trait in question.[20] Instead, the explanation of that fact depends on other processes. Novel phenotypic traits arise in a population on the basis of genetic mutation and recombination. These mutations are due to stray cosmic rays, errors in replication, and other chance events. These traits are then passed on from organisms to their descendants.[21] The explanation of how a particular organism came to have a phenotypic trait thus does not involve selection but heredity and extremely chancy events.

It is important to be careful here. Natural selection *can* help to explain why a novel phenotypic trait arose in the population. Previous bouts of selection could have made it very likely that the trait would emerge.[22] But what selection cannot do is explain why a particular individual has the trait in question.

This feature of selective explanations might provoke the following worry: To fully respond to the reliability challenge, one must provide a satisfying explanation both of why the *population* primarily includes thinkers who employ reliable deductive rules

[20] See Sober (1984), section 5.2.
[21] I am here ignoring the complexities that arise from the impact of the environment on phenotypic traits and from the existence of sexual reproduction. These complexities do not affect the points to follow.
[22] See Neander (1995). Neander also claims that selection can explain why a particular organism has a phenotypic trait. I don't see why this further claim is true.

and of why *particular thinkers* employ reliable deductive rules. Whatever the prospects are for an evolutionary account to meet the first demand, it is unable to meet the second. And it would be highly unsatisfying to be forced to think that it was merely an accident that each individual thinker employs reliable deductive rules. The claim that it was merely an accident would produce a tension in our overall view of the world.

While this worry is arresting, it is ultimately misguided. The primary fact to be explained is a population-level fact – our population primarily includes thinkers that employ reliable deductive rules. To illustrate this point, suppose that there were a vast population of heterogeneous thinkers, each with a different set of inferential rules. Suppose that a few of these thinkers had reliable deductive rules, and that the number of such thinkers was roughly what one would expect if the rules were somehow distributed randomly. If the few 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 they were reliable. (Compare: Given that we have many sets of dice, it is not striking that some particular set came up all sixes. What is striking is that every set of dice in our possession came up all sixes.) This suggests that the reliability of a particular individual is not in general a striking fact. What is striking is the reliability of the population at large.

There is a complication, however. Consider the first-personal claim that *I* employ reliable deductive rules. It is intuitive that this is a striking fact. At the very least, this fact seems more in need of explanation than the claim that some arbitrary particular individual is reliable. When I reflect on my own reliability, it is hard to be satisfied with the thought that it was merely a lucky fluke that I am reliable. Since an evolutionary

account cannot explain this first-personal fact, we again have a worry about the prospects of an evolutionary account.

I'm not entirely sure what to make of this concern. There is reason to be cautious here. It is not entirely clear that we should endorse the intuition that it is a striking fact that I am reliable. It is not obvious that my winning the lottery would be any more striking than the fact that some particular person won the lottery. And even if the first-personal fact is a striking fact, it might not be a big cost of a view to claim that this fact came about merely by accident. But it should be granted that there are murky waters here.[23]

Putting the first-personal case aside, assuming that an evolutionary account of the reliability of our population can be made to work, the fact that it is accidental that particular thinkers employ reliable rules is not worrisome. But can an evolutionary account of the reliability of our population be made to work?

4. An Evolutionary Explanation of Our Deductive Reliability

Deductive rules are rules of inference that govern reasoning with logical concepts. One natural strategy for explaining why it is that we employ reliable deductive rules is to divide the task into two parts. The first is to explain why it is that we possess logical concepts. The second is to explain, given that we possess logical concepts, why it is that we employ reliable deductive rules. In this section, I sketch an evolutionary account with this bipartite structure. According to this account, possessing logical concepts and

[23] In conversation, David Christensen has suggested that although selection cannot explain the first-personal fact that I have a reliable deductive mechanism, it nevertheless raises its probability. While this seems true, it does not assuage the worry. The general principle of theory choice generating the reliability challenge concerns explanation and not probability.

employing reliable deductive rules both conferred survival advantages on our ancestors. This explains why it is that we have these phenotypic traits.

4.1. The Advantages of Possessing Logical Concepts

On the proposal under consideration, possessing logical concepts conferred a survival or reproductive advantage on our ancestors. But it is *prima facie* puzzling what this advantage could be. The main difficulty is as follows: It is tempting to think that the inputs from the world (via perception) to our reasoning faculties are logically simple. It is also tempting to think that our behavior and behavioral dispositions depend only on the logically simple products of our reasoning. Given these two assumptions, it is plausible that being able to hold logically complex beliefs can confer a selective advantage only if such beliefs are important as intermediate steps in our reasoning. But this is apparently ruled out by the fact that logical concepts are *conservative* in the following sense:[24] Adding logical concepts (along with the associated deductive rules) to any conceptual practice involving only logically simple propositions does not license any inferences from logically simple premises to logically simple conclusions that were not already licensed before the addition.[25]

In response, it might be suggested that possessing conservative concepts can be advantageous in another way – namely, by enabling thinkers to reason in a quicker or

[24] Conservativeness (or a related property) is often taken to be a necessary condition on logical constanthood. This proposal stems from the Gentzen-Prawitz tradition of requiring the introduction and elimination rules of a logical concept to appropriately match each other, and from the suggestion in Belnap (1961) that a conceptual role bestows a genuine meaning on a logical constant only if it satisfies a conservativeness requirement. Hacking (1979) demarcates the logical, in part, by appealing to conservativeness. See Dummett (1991) and Tennant (1987) for relevant discussion.

[25] In the absence of a detailed list of the deductive rules we employ, a rigorous proof of conservativeness is not possible. But this claim, or a related one, is undoubtedly correct.

more efficient manner than they would otherwise be able to.[26] But this suggestion does not help for the case of logical concepts. Adding logical concepts (with the associated deductive rules) to a conceptual practice does not speed up derivations of logically simple conclusions from logically simple premises.

In answer to this difficulty, I suggest that the two tempting thoughts ought to be rejected. Logically complex beliefs can be inputs to reasoning, delivered by perception or other non-inferential cognitive mechanisms. This is perhaps clearest for conjunctive propositions. Our visual system binds together different features of an object and delivers beliefs in conjunctions, such as the proposition that a given object is both round and red. Beliefs delivered by perception may be logically complex in other ways, too. Consider, for instance, negative existential propositions. It seems that just by looking – without any inferential reasoning – I can come to believe that there is no dog in my office.[27] Similarly, perhaps thinkers can directly come to believe disjunctive propositions on the basis of perception, such as the proposition that a certain object is either blue or green.[28] Perhaps, too, thinkers can directly come to believe general propositions, such as the proposition that every book on the top shelf is blue. Moreover, if beliefs acquired through testimony are not based on inference, arbitrary logically complex beliefs may be inputs to reasoning. On the basis of logically complex inputs to reasoning, thinkers may reason to logically simple conclusions and behave accordingly.

---

[26] Compare: Mathematical theories may be useful when added to nominalist physical theories by shortening the derivation of purely nominalist physical consequences from nominalist physical claims. See Field (1980).

[27] It might be claimed that my belief that there is no dog in my office is not solely based on perception, but also on a default presumption that what I don't see isn't there. This view faces several difficulties. But even if some version of it is correct, the default presumption is itself a logically complex input to reasoning.

[28] See Dummett (1991), page 267, for the claim that perceptions can be disjunctive. Dummett's example is that Hardy may not have been able to tell whether Nelson said, "Kismet, Hardy" or "Kiss me, Hardy", but may have had a perception with an "irreducibly disjunctive form".

Logically complex beliefs can also be important in guiding behavior, and not merely by virtue of their logically simple consequences. For a mundane example, suppose that I am searching for a glass of water. Holding the disjunctive belief that a glass contains either water or vodka should have a different impact on my behavior than a belief in either of the disjuncts. If I were to believe that the glass contains water, I should pick up the glass and abandon my search. If I were to believe that it contains vodka, I should ignore it and continue my search. But if I were to believe the disjunction, I should examine the liquid more closely, perhaps take a sip, and perform other appropriate tests.[29] Similarly, holding the negative belief that there is no glass of water nearby should motivate me to search elsewhere. Holding the conditional belief that if there is a glass of water nearby then it is in my office should motivate me to look in that location. In general, possessing logical concepts enables thinkers to represent important information about the world, information of use in guiding behavior.

There is a natural picture of the representational abilities provided by the logical concepts. Consider the space of all (metaphysically or epistemically) possible worlds. We can think of each logically simple proposition as picking out a region of this space – namely, the worlds at which the proposition is true. The logical concepts enable us to pick out additional regions of this space in our thought. Conjunction enables us to take the intersection of regions. Disjunction enables us to take the union of regions. Negation enables us to take the complement of a region. And so on. Being able to pick out these

[29] For an analogous example, see Skyrms (1999). In discussing a hypothetical group of logically sophisticated vervet monkeys, Skyrms claims that the optimal evasive action given knowledge that there is either a snake or a leopard nearby may be different both from the optimal action given the knowledge that there is a leopard nearby and from the optimal action given the knowledge that there is a snake nearby.

additional regions of modal space helps us to represent information that we acquire from the world and make use of in guiding our behavior.

As we will see below, this picture is much too crude. But it provides a useful initial account of the role of the logical concepts in thought.

It is interesting to note that it is not necessary for a thinker to possess logical concepts in order to have the representational abilities they provide. For instance, instead of believing of something that it is either blue or green, a thinker could believe that it is blue-or-green, where *blue-or-green* is a concept that applies to blue things and to green things. Instead of believing that it is not the case there is a chair here, a thinker could believe that it is un-chaired here, where *un-chaired* is a concept that applies to locations lacking chairs. Similar claims hold for more complex logical constructions. Instead of possessing a few general purpose logical concepts, thinkers could possess a bevy of special purpose concepts. However, this does not provide an objection to the claim that possessing logical concepts can be evolutionarily advantageous. It would be computationally very costly for a thinker to dispense with general-purpose logical concepts in favor of a large number of special purpose concepts. Moreover, to have the representational abilities that logical concepts provide, such a thinker would have to add a potentially unbounded number of additional concepts whenever a new concept was acquired.

There is an additional proposal about the advantage of possessing logical concepts that merits discussion. This is the idea (inspired by the work of Brandom) that the role of logical concepts is to enable thinkers to “make explicit their implicit inferential

commitments."[30] The idea here is that thinkers are disposed or committed to infer in certain ways. Logical concepts provide thinkers with the conceptual resources needed for subjecting these commitments to rational scrutiny. Consider, for instance, *the conditional*. In believing *if A then B*, a thinker makes explicit her commitment to inferring B from A.[31] This belief can then be assessed by considering reasons for and against it. If the latter are more persuasive than the former, the belief – and the inference – may be rejected. Adapting this view to an evolutionary context, the idea is that possessing logical concepts such as *the conditional* conferred a survival advantage because they enabled our ancestors to hold up their inferential commitments to rational scrutiny, and thus enabled them to improve their methods of reasoning about the world.[32]

This suggestion is appealing. It is plausible that an important role of logical concepts is to allow thinkers to better assess their inferential practices and commitments. However, the proposal need not be seen as a competitor to the idea that the function of the logical concepts is to enable thinkers to represent facts about the world that they

---

[30] See Brandom (2000, introduction and chapter 1). Brandom's view concerns logical vocabulary rather than logical concepts.

[31] This is reminiscent of the view in Ryle (1950) that conditionals are inference tickets. Brandom sometimes writes that the job of conditionals is to enable speakers to *say* that certain inferences are acceptable. Presumably, this should be construed in a loose sense. "If A then B" does not say anything about an inference (provided that neither A nor B do so). Rather, it says something about the world. The best way to understand Brandom's view, I take it, is that in stating "if A then B", a speaker asserts something about A and B, and in so doing also *expresses* a commitment to the correctness of the inference from A to B.

[32] Consider the following conditional: If Sally is deceiving me, I'll never believe that she is (because she is so clever). Presumably, a thinker can sensibly believe this conditional without being committed to reason from the antecedent to the consequent. So there seems to be a counterexample to this view of conditionals. (This kind of example is originally due to Thomason. See van Fraassen (1980), page 503, for the attribution.) There are several things that can be said in reply. One could claim that a thinker who endorses the conditional does express a commitment to make the corresponding inference but loses this commitment upon coming to believe the antecedent. Alternatively, one could modify the view by claiming that in endorsing a conditional, a thinker expresses a commitment to viewing the corresponding inference as a good inference (but does not necessarily express a commitment to drawing the inference). Finally, it could be claimed that the semantics of conditionals and the selective advantage of possessing *the conditional* come apart. Although endorsing a conditional does not always go along with a commitment to reasoning from the antecedent to the consequent, it typically does. This is enough to explain why possessing *the conditional* conferred a survival advantage. Thanks to an anonymous referee for pressing me on this issue.

would not otherwise easily be able to represent. Possessing logical concepts may have conferred a survival advantage *both* because they improved our ancestors' representational abilities and because they helped our ancestors better assess their own patterns of reasoning. There is no need to choose between these two suggestions.[33]

4.2. The Advantages of Employing Reliable Deductive Rules

We see, then, that possessing logical concepts can confer an evolutionary advantage. What about the employment of reliable deductive rules? This question raises an old problem about deduction. The problem is to explain the point of deductive inference.[34] Deductive inference does not augment our knowledge of the world. In an intuitive sense, the conclusion of a deductively valid inference does not contain any information not already contained in the premises. This claim has a formal analogue: The set of possible worlds at which the conclusion is true is a superset of the possible worlds at which all of the premises are true. What, then, could the purpose of deductive inference possibly be? How could there be an evolutionary advantage in employing reliable deductive rules?

One strategy for answering this question is to appeal to some version of Conceptual Role Semantics. On this view, for a thinker to possess a concept, she must employ certain associated rules of inference or belief-forming methods. Assuming that Conceptual Role Semantics holds true for logical concepts, it is plausible that the deductive rules we employ are concept-constituting of the logical concepts we possess.

[33] Here are two additional proposals: First, possessing logical concepts is needed to engage in certain kinds of non-deductive reasoning, such as inference to the best explanation. Second, possessing logical concepts such as *the conditional* enables thinkers to appropriately hedge their beliefs and avoid unnecessary error. The latter suggestion is related to the discussion of *the conditional* in Boghossian (2003).
[34] This problem goes back at least as far as Mill (1843), II.3.1.

So any advantages there are for possessing logical concepts are also advantages for employing the associated deductive rules.

Although this strategy has some appeal, there are reasons for doubt. First, Conceptual Role Semantics is subject to many difficulties, and may well be false.[35] Second, the explanation is intuitively the "wrong way around". On many versions of Conceptual Role Semantics, possessing a concept *consists in* employing certain rules. So one would expect the evolutionary advantages of possessing a concept to derive from the advantages of employing its associated rules, and not vice-versa. Finally, and more importantly, the strategy is at best able to explain how it is we employ *particular* deductive rules – those constitutive of the particular logical concepts we possess. It is unable to explain how it is we employ *reliable* deductive rules. The reliability of the rules plays no obvious role in the explanation.[36] The explanation makes it seem to be merely an accident that we came to employ reliable deductive rules.[37]

There is a better strategy for explaining the evolutionary advantage of employing reliable deductive rules. A preliminary point to make is that information from the world does not come to us all at once in a single package. We acquire it over time and from disparate sources. Deductive reasoning can help us combine new information with old information. It can help us combine information from different sources. It can also help us to reassess old information or to apply old information to new situations.

[35] For some of the difficulties, see Block (2000), Fodor and Lepore (1991), and Williamson (2003).
[36] To be fair, on some versions of Conceptual Role Semantics, constitutive rules are required to be truth-preserving. For instance, see Peacocke (1992). However, there are compelling counterexamples to this claim. See Boghossian (2003) and Schechter and Enoch (2006) for discussion.
[37] A more plausible suggestion is that the selective advantage of possessing logical concepts required their constitutive rules to be reliable. This suggestion is compatible with the one to follow.

Yet, simply saying this does not suffice to explain the evolutionary advantage of employing reliable deductive rules. When a thinker employs a reliable deductive rule to combine bodies of information or to draw out a consequence of some body of information, the thinker is not acquiring any novel information. All of the relevant information is already present, at least on a coarse-grained understanding of "information". So what could the evolutionary advantage of employing a deductive rule be?

The answer to this question must be this: Employing reliable deductive rules helps thinkers to convert information into a more usable form. To illustrate, suppose again that I am searching for a glass of water. Even if I were to come to believe both that there is a glass of water located either here or over there and that the glass of water is not located here, I would not immediately be motivated to look for it over there. In order for me to be so motivated, I would first have to draw the relevant inference. I would have to come to explicitly believe that the glass of water is over there. This is so despite the fact that the information that the water is over there is implicit in my original beliefs. (This point is still more striking for cases where the relevant sequence of deductive inferences is very long.) In general, rational behavior depends not only on the information a thinker possesses, but also the manner in which it is represented.

This shows that the natural picture of the role of the logical concepts in thought presented above – in terms of picking out regions of modal space – is too crude. Logically equivalent propositions can be importantly different in their relevance to behavior. A deductive consequence of some premise may pick out a larger region of modal space than the premise, but may also be directly relevant to an urgent matter in a

way that the premise was not. Part of the role of the logical concepts, then, is to represent information in specific ways. Deductive rules of inference are important because they enable us to manipulate the ways that information is represented.

The proposal, then, is this: Employing reliable deductive rules enables thinkers to better make use of the information they possess in guiding their behavior. This is why employing reliable deductive rules conferred an evolutionary advantage on our ancestors.

5. Objections and Replies

We have, then, a (very rough) sketch of an evolutionary explanation of our reliability about logic. We possess logical concepts and employ reliable deductive rules because this conferred an evolutionary advantage on our ancestors. Possessing logical concepts enabled our ancestors to represent important information about the world as well as to assess their own patterns of reasoning. Employing reliable deductive rules enabled our ancestors to better make use of the information in their possession. Our reliability about logic is a side-effect of our employment of reliable deductive rules.

It may be worthwhile to here repeat an earlier point: My discussion is only intended to provide a sketch of an account. To fill out the explanation would require a considerable amount of empirical work. For instance, one might want to look into the following issues: How is it that our cognitive mechanisms are genetically encoded? How is it that natural selection works on the mind? What are the intermediate forms between animals not possessing cognitive mechanisms for deductive reasoning and animals that do possess such mechanisms? Why are humans special in having such powerful cognitive capacities? What is the specific role of each individual logical constant and each

individual deductive rule in our thought? What are the costs of having a cognitive mechanism for deductive reasoning? And so on.

It is not one of the tasks of this paper to answer these interesting and important empirical questions. The central goal of this paper, rather, is to address the reliability challenge for logic and to thereby ward off the pressure on our belief that logic is objective, on our belief that we are reliable, and on our general background views about the world. This is a task in the epistemology of logic. Providing a plausible sketch of an explanation, even at a high level of generality, is sufficient for achieving this goal.

There are, however, several pressing objections that can be raised against evolutionary accounts of our reliability about logic. These objections purport to show that such accounts cannot even in principle explain our reliability about logic. In this section, I discuss the most serious of these objections. I show how they may best be addressed. The objections fall into two categories. The first concerns whether evolutionary accounts can explain the full extent of our reliability. The second concerns whether evolutionary accounts can explain our reliability at all. In responding to these objections, I further develop the evolutionary account.

5.1. Generality, Complexity, and Necessity

One difficulty for evolutionary explanations of our reliability is that they seem unable to explain the full extent of our reliability. It is plausible that our ancestors were selected for correctly reasoning only about certain domains. Presumably, what natural selection "cares about" is that our ancestors formed true beliefs about danger, food, mating, shelter, and other topics closely tied to survival and reproduction. Perhaps when sexual selection

is taken into account, the list can be somewhat expanded. But even if it can, it is clear that our ancestors were not selected for correctly reasoning about algebraic geometry, the Problem of the Many, the geology of Mars, or how best to make a cheese soufflé. Evolutionary accounts thus seem unable to explain how it is that we employ deductive rules that are reliable for arbitrary subject matters.[38]

A second difficulty concerns the complexity of propositions. It is plausible that our ancestors were selected for reasoning correctly only with relatively simple propositions. Presumably, there was little survival advantage in reasoning correctly with highly complex propositions. Evolutionary accounts thus seem unable to explain how it is that we employ deductive rules that are reliable when applied to propositions of arbitrary complexity.

A third difficulty comes from the reaches of modality. The reliability of our deductive mechanism is not merely a matter of its yielding true beliefs from true beliefs. Our deductive mechanism is reliable in a modally robust sense – the rules we employ *necessarily* preserve truth. Evolutionary accounts seem unable to explain this fact. Success in survival and reproduction depends only on reasoning in ways that are truth-preserving in the actual world, and not on reasoning in ways that would be truth-preserving even in far off possible worlds.[39]

---

[38]A related worry appears in Delbrück (1978), page 353, and is discussed in Sober (1981). Delbrück's worry concerns the mechanisms used in scientific reasoning. Nagel (1986), page 79, argues against evolutionary explanations of our capacity for objective thought on the grounds that "our capacity to form cosmological and subatomic theories takes us so far from the circumstances in which our ability to think would have had to pass its evolutionary tests…." Field (1989), pages 28–29, argues against evolutionary explanations of our reliability about mathematics on the grounds that only a small portion of mathematics is tested against the world.

[39] Nozick (2001), page 122, and Stroud (1981) argue that evolutionary accounts make it difficult to see how we could be reliable in our attributions of necessity, since selection only rewards believing truths about the actual world. This objection is different from the one I discuss. I do not here consider the status of attributions of modal properties.

Indeed, when carefully evaluated, it seems incoherent to claim that our ancestors were selected for employing deductive rules that are necessarily truth-preserving. A necessary condition for there to have been selection for a trait is that the presence of the trait *caused* the relevant organisms to survive longer or reproduce more frequently. How could it be that a modal property was causally efficacious in this way?

In response to the third of these difficulties, one might try to deny that our deductive rules are reliable in a modally robust sense. Perhaps the deductive rules we employ would be truth-preserving if the world were only slightly different, but would not be truth-preserving if the world were radically different.[40] However, this response is unattractive. Adopting it would be tantamount to giving up on many of our modal beliefs: If we were to believe that even our deductive rules are not necessarily truth-preserving, we should believe that we are not at all reliable about what is necessarily the case.

Moreover, an analogous response cannot be used to answer the first difficulty. It would be extremely difficult to maintain that our deductive rules produce true beliefs when reasoning about predators but not when reasoning about other topics, such as how to get an astronaut to the Moon. Such a view conflicts with the success we have had in many of our endeavors. Such a view is also potentially self-undermining: Any belief that we are unreliable concerning subject matters that were not of immediate evolutionary relevance should undermine our trust in the very reasoning that produced the belief.

There is no plausible way to argue that our ancestors were selected for reasoning correctly about arbitrary subject matters, with propositions of arbitrary complexity, and in

[40] A related proposal is put forward by Nozick (1993), page 111, who suggests that the explanation of why logical principles seem self-evident may be that "they are true, even if only contingently, even just 'true enough' … and that they have held true durably for long enough to leave an imprint upon our evolutionary endowment." This is reminiscent of the view in Mill (1843), II.6.2, that arithmetic is contingent but appears necessary due to our "early and constant experience".

a way that works in arbitrary possible worlds. The answer to the three difficulties must instead be that these features of our deductive rules are *byproducts* of the traits for which our ancestors were really selected. Our ancestors were selected for employing deductive rules that actually are truth-preserving for a limited range of propositions. Our ancestors were also selected for employing rules that are computationally efficient, have few storage costs, and the like. A good way to satisfy these constraints is to employ deductive rules that work generally and necessarily, and this explains why our ancestors came to have such rules.

This response might be thought to be somewhat disappointing. What we were after was a satisfying explanation of how we came to have generally and necessarily truth-preserving deductive rules, and the evolutionary explanation on offer claims only that this is a side-effect of something else. But this disappointment should be fleeting. There is nothing inherently problematic about the style of explanation offered. Indeed, evolutionary biologists are familiar with numerous examples in which the most plausible explanation of some trait is that it is a side-effect of what was really selected for.[41]

What would be unsatisfying is if it were claimed that our employment of reliable deductive rules was a side-effect of a trait that was not directly related to reliability, such as the trait of having large ears. This would make it seem accidental that we came to employ deductive rules that work generally and necessarily. But the explanation under consideration makes no such claim. The trait of employing rules of inference that are truth-preserving in a limited range of cases is closely tied to the trait of employing rules

[41] Gould and Lewontin (1978) use the word "spandrels" to refer to traits that are not selected for but are side-effects of how an organism develops or is built.

that generally and necessarily preserve truth. The proposal thus provides a satisfying answer to the three difficulties.

5.2. USEFULNESS AND TRUTH

The final difficulty for evolutionary accounts of our reliability is the most general. Namely, even if an evolutionary account can explain why we have useful cognitive mechanisms, it is not clear it can explain why we have cognitive mechanisms that are truth-conducive.

A version of this objection has been forcefully presented by Stich.[42] Stich directly argues only that natural selection does not *guarantee* that our cognitive mechanisms generally produce true beliefs (given the appropriate inputs). This conclusion is unobjectionable. But if his arguments can be extended to support the claim that natural selection is *unlikely* to yield cognitive mechanisms that generally produce true beliefs, this would count against an evolutionary explanation of our reliability.

Stich argues for two claims. First, evolution does not always yield the optimal mechanism for a task. That is, it does not always produce a mechanism that provides the greatest contribution to fitness when compared to alternatives. Second, optimal cognitive mechanisms are not always the most reliable. Features other than truth-conduciveness may be more important to fitness.

In support of the first claim, Stich presents several reasons why evolution does not always yield an optimal design: An optimal mechanism may not be biologically possible.

[42] See Stich (1990, chapter 3). Stich does not challenge the view that our cognitive mechanisms are the products of evolution, but rather the view that "evolutionary considerations impose interesting limitations on irrationality." Churchland (1987) argues that there is a significant gap between behaving in a way conducive to survival and having mostly true beliefs.

It may not be biologically possible for a population to evolve an optimal mechanism given its current characteristics. Random genetic drift may lead to the disappearance of an optimal mechanism from a population or may hamper its spread. The complexities of sexual reproduction may do so, as well.

Stich is undoubtedly correct that evolution does not always maximize fitness. However, Stich provides no reason to believe that evolution is unlikely to provide mechanisms that substantially enhance fitness. Moreover, it is difficult to evaluate how well Stich's considerations apply in the particular case of interest. Not much is known about the evolution of the brain, and more specifically, the evolution of cognitive traits. It is an open question whether any of the factors Stich discusses were of any real importance in the evolution of our reasoning mechanisms.

In support of the second claim, Stich argues that there is frequently a tradeoff between how truth-conducive a cognitive mechanism is and how economical the mechanism is with respect to the amount of energy, time, and cognitive hardware that it requires. It may be more important to an organism's survival that a cognitive mechanism be “fast and frugal” than that it be generally accurate. In addition, there may be a tradeoff between the number of false positives and the number of false negatives that a cognitive mechanism produces. It may be very important to avoid one of these – for instance, false negatives concerning the existence of lurking predators – at the cost of substantially increasing the other.[43]

Again, it is undoubtedly correct that the fittest cognitive mechanism is not always the most truth-conducive. Yet, there are general reasons to think that the gap between the two is not as large as Stich's discussion might lead one to expect. All things being equal,

[43] As Stich notes, these points are originally due to Sober (1981).

correctly representing the world is conducive to survival; it helps organisms react to threats and achieve important goals.[44] It seems likely that the cognitive mechanisms that evolution produces are at least somewhat truth-conducive in a central range of cases.

There is also reason to believe that Stich's considerations carry less force for the particular case of deductive reasoning. The structure of logical entailments is relatively simple. Indeed, cooking up rules that resemble the deductive rules we employ but which are not universally or necessarily truth-preserving would seem to require adding *ad hoc* conditions or otherwise increasing their complexity. So for the specific case of deductive reasoning, there seem to be no additional computational costs in employing fully truth-conducive rules over somewhat-but-not-fully-truth-conducive rules. There may be no less reliable but more economical cognitive mechanism in the ballpark. Moreover, deductive reasoning is a general-purpose reasoning mechanism. So there is little pressure for it to minimize false positives at the expense of false negatives, or vice-versa. There is no obvious reason for selection to yield any systematic bias for one over the other.

There is an additional point to make in response to Stich's claim. Namely, there is reason to think that that there was in fact selection pressure for the employment of deductive rules that are highly truth-conducive in an important range of cases.

In general, it can be advantageous for organisms to have cognitive mechanisms that are highly truth-conducive albeit slow or inefficient. In certain circumstances – where energy and time is available, the stakes are high, and the relevant issues are tricky – it can be important not to form beliefs rashly, but rather to have a more considered

[44] Stich (1990, chapter 5), argues that it is not at all valuable to have true beliefs. His argument depends on the difficulty of assigning truth conditions to mental states in a privileged way. This is a deep issue in the philosophy of mind and cannot be addressed here.

response. This is a point familiar from everyday life, and it extends to the context of evolutionary fitness.

Deductive reasoning has the hallmarks one would expect of a cognitive mechanism that was selected for playing this role. Deductive reasoning is relatively slow, at least compared to instinctive responses or the application of simple heuristics. Its outputs are only indirectly tied to behavior. It is content-neutral, capable of representing and transitioning between a large range of contents. And it is highly reliable. What these features suggest is that our deductive mechanism was designed to get it right when getting it right is very important and time is not of the essence.[45]

This suggestion fits well with "dual process" theories in psychology, according to which we possess two cognitive systems.[46] System 1 is a relatively fast, cognitively undemanding, automatic reasoning system that makes use of special-purpose heuristics. System 2 is a relatively slow, cognitively demanding, general-purpose system for explicit reasoning. Our mechanism for deductive reasoning is a part of System 2. This system is highly reliable, at least in an important range of cases.

The claim, then, is that one of the functions of our deductive mechanism requires it to be highly truth-conductive for some range of cases. This raises the question: What is that function? Above, I proposed that the evolutionary advantage of employing deductive rules is that they enable thinkers to convert information into a form that makes it available for guiding behavior. This goes some way towards answering the question. But

[45] Papineau (2000) argues for an analogous view about a different cognitive mechanism. He claims that we have a specialized means-end reasoning module that selects the best course of action in the light of the available information when "the stakes are high and time does not press."
[46] See Evans and Over (1996) and Sloman (1996).

the question remains: For what purpose (or purposes) did our ancestors need to represent logically complex information and manipulate it in a highly truth-conducive way?

There is one suggestion that I find both plausible and illuminating. Namely, employing highly truth-conducive deductive rules conferred a selective advantage on our ancestors because it helped them to successfully engage in long-term planning.[47]

There are two types of long-term planning. The first is what might be called "strategic planning". This is the sort of planning a thinker engages in when constructing a strategy for governing her actions over some stretch of time in order to achieve some desired end. The second is what might be called "contingency planning". This is the sort of planning a thinker engages in when constructing a plan for handling potential future emergencies or opportunities.

Three general features of long-term planning suggest a connection with our employment of reliable deductive rules. First, long-term planning typically relies on many different pieces of information, provided by many different cognitive mechanisms. This suggests that deductive reasoning plays an important role in long-term planning; such reasoning enables thinkers to combine the relevant information and transform it into a more usable form.

Second, long-term planning is an important cognitive endeavor, one in which the benefits of success and the costs of failure can be high. We do not always have the time or energy to deliberate about what to do at the time we need to act; long-term planning enables us to ameliorate the impact of such resource constraints. Such planning also enables us to better coordinate our actions (with ourselves and with others); it increases

[47] I originally owe this suggestion to Paul Boghossian.

the likelihood that we will achieve important ends.[48] This suggests that success in long-term planning provides a significant evolutionary advantage.

Finally, when engaging in long-term planning, there is typically little need to make an immediate decision. This suggests that there is little selection pressure against the employment of relatively slow cognitive mechanisms in planning, at least so long as there are compensatory benefits.

Taken together, these three features suggest that there was a strong selective pressure for the employment of highly truth-conducive cognitive mechanisms to use in such planning, even if such mechanisms are slower or less efficient than alternatives. The proposal, then, is that our ancestors were selected for employing highly truth-conducive deductive rules in part because it helped them engage in long-term planning.[49]

Other proposals can also be envisioned. For instance, a currently fashionable view is that many of our cognitive abilities were developed to help with the complexities of living in social groups. On one version of this proposal, our cognitive abilities were developed to help to help compete with conspecifics for limited resources.[50] On a different version of this view, our cognitive abilities were developed to help cooperate with conspecifics in engaging in complicated cooperative endeavors such as foraging.[51] Alternatively, it could be that employing highly truth-conducive rules was important for the task of constructing explanations of important phenomena – such as the causes of certain animal traces.[52] A still different proposal is that our cognitive abilities are the product of sexual selection. Being able to reason clearly and in a truth-conducive way

[48] These claims about the importance of planning are due to Bratman (1987; 2000).
[49] Similar remarks can be made about deliberation more generally.
[50] See, for instance, Humphrey (1976) and Flinn, Geary, and Ward (2005).
[51] See Sterelny (2003; 2007).
[52] Carruthers (2002b) suggests that explanatory reasoning originally emerged to help hunters track animals.

were indicators of the fitness or an organism, and so organisms that were more intelligent were favored as mates.[53] There are many other alternatives, too. Many of these views have the resources to explain why we came to employ highly truth-conducive cognitive mechanisms for deductive reasoning.

The general moral to draw is not that any particular view (or views) ought to be adopted. Rather, it is that natural selection can sometimes yield highly truth-conducive cognitive mechanisms. Our deductive mechanism is a prime candidate for being such a mechanism.

6. Conclusion

On the proposed account, we are reliable in believing logical truths and disbelieving logical falsehoods because we have a reliable mechanism for deductive reasoning. The explanation of how we came to have a reliable deductive mechanism is an evolutionary one. Our ancestors were selected for possessing logical concepts. Possessing these concepts conferred an evolutionary advantage because it enabled our ancestors to represent important information about the world and because it aided them in assessing their own patterns of reasoning. Our ancestors were also selected for employing deductive rules that are truth-preserving in an important range of cases. Employing such rules enabled our ancestors to convert information into a form that was more useful for guiding their behavior. (According to one proposal, it was advantageous for our ancestors to do this in a highly truth-preserving way because of the importance of long-term planning.) There was an additional selection pressure on our ancestors, namely, to

[53] See Miller (2000).

employ rules that are cognitively economical. As a byproduct of these pressures, our ancestors came to employ rules that are generally and necessarily truth-preserving. Employing reliable deductive rules is a heritable trait, and so we came to employ them, too.

On this proposal, our reliability about logic is a byproduct of a byproduct. It is a byproduct of the general reliability of our cognitive mechanism for deductive reasoning. This in turn is a byproduct of what was really selected for – an economical cognitive mechanism that is highly reliable in a central range of cases.

This proposal provides a satisfying answer to the reliability challenge for logic. It demonstrates that there are plausible answers to be had. Whether or not the precise details of the account are correct, it shows that there is no in principle reason to think that evolutionary accounts are incapable of explaining our reliability about logic. This defuses the tension between the claim that logic is objective, the claim that we are reliable about logic, and our general background views about the world.[54]

[54] A much earlier version of this paper was presented at an NYU dissertation seminar. I would like to thank the participants for their questions and remarks. I would also like to thank Paul Boghossian, Ray Buchanan, Winston Chiong, David Christensen, Jamie Dreier, David Enoch, Greg Epstein, Dana Evan, Hartry Field, Kit Fine, Don Garrett, Pete Graham, Liz Harman, Jonathan Jenkins Ichikawa, Øystein Linnebo, Anna-Sara Malmgren, Gary Marcus, Christopher Peacocke, Erica Roedder, Katherine Rubin, Katia Samoilova, Karl Schafer, Brad Skow, Declan Smithies, Sharon Street, Michael Strevens, Roger White, and Masahiro Yamada for valuable discussion. I would also like to thank two anonymous referees for their helpful comments.

REFERENCES

Belnap, N. (1962) 'Tonk, Plonk and Plink', *Analysis*, 22: 130–134.

Benacerraf, P. (1973) 'Mathematical Truth', *The Journal of Philosophy*, 70: 661–679.

Block, N. (2000) 'Semantics, Conceptual Role', in E. Craig (ed.), *Routledge Encyclopedia of Philosophy* (London: Routledge).

Boghossian, P. (2003) 'Blind Reasoning', *Proceedings of the Aristotelian Society, Supplementary Volume*, 77: 225–248.

Braine, M., and D. O'Brien, eds. (1998) *Mental Logic* (Mahweh: Erlbaum).

Brandom, R. (2000) *Articulating Reasons: An Introduction to Inferentialism* (Cambridge, Mass.: Harvard University Press).

Bratman, M. (1987) *Intention, Plans, and Practical Reason* (Cambridge, Mass.: Harvard University Press).

——— (2000) 'Reflection, Planning, and Temporally Extended Agency', *Philosophical Review*, 109: 35–61.

Carruthers, P. (2002a) 'The Cognitive Functions of Language', *Behavioral and Brain Science* 25: 657–674.

——— (2002b) 'The Roots of Scientific Reasoning: Infancy, Modularity, and the Art of Tracking', in P. Carruthers, S. Stich, and M. Siegal (eds.), *The Cognitive Basis of* Science (Cambridge: Cambridge University Press), 73–98.

Churchland, P. (1987) 'Epistemology in the Age of Neuroscience', *The Journal of Philosophy*, 84: 544–553.

Darwin, C. (1871) *The Descent of Man and Selection in Relation to Sex* (New York: D. Appleton and Company).

Delbrück, M. (1978) 'Mind from Matter?', *American Scholar*, 47: 339–353.

Dummett, M. (1991) *The Logical Basis of Metaphysics* (Cambridge, Mass.: Harvard University Press).

Evans, J. and D. Over (1996) *Rationality and Reasoning* (East Sussex: Psychology Press).

Field, H. (1980) *Science without Numbers: A Defense of Nominalism* (Princeton: Princeton University Press).

——— (1989) *Realism, Mathematics and Modality* (Oxford: Basil Blackwell).

——— (1998) 'Epistemological Nonfactualism and the A Prioricity of Logic', *Philosophical Studies*, 92: 1–24.

Flinn, M., D. Geary, and C. Ward (2005) 'Ecological Dominance, Social Competition and Coalitionary Arms Races', *Evolution and Human Behavior*, 26: 10–46.

Fodor, J. and E. Lepore (1991) 'Why Meaning (Probably) Isn't Conceptual Role', *Mind and Language*, 6: 328–34.

van Fraassen, B. (1980) 'Review of *Rational Belief Systems* by Brian Ellis', Canadian *Journal of Philosophy*, 10: 497–511.

Gould, S. and R. Lewontin (1978) 'The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme', *Proceedings of the Royal Society of London*, 205: 581–598.

Hacking, I. (1989) 'What is Logic?', *The Journal of Philosophy*, 76: 285–319.

Harman, G. (1988) *Change in View: Principles of Reasoning* (Cambridge, Mass.: The MIT Press).

——— (1995) ‘Rationality’, in E. Smith and D. Osherson (eds.), *Thinking: An Invitation to Cognitive Science*, volume 3, second edition (Cambridge, Mass.: The MIT Press), 175–211.
Hill, C. (2006) ‘Modality, Modal Epistemology, and the Metaphysics of Consciousness’, in S. Nichols (ed.), *The Architecture of the Imagination: New Essays on Pretense, Possibility, and Fiction* (Oxford: Oxford University Press), 205–235.
Humphrey, N. (1976) ‘The Social Function of Intellect’, in P. Bateson and R. Hinde (eds.), *Growing Points in Ethology* (Cambridge, UK: Cambridge University Press), 303–317.
Mill, J. S. (1843) *A System of Logic, Ratiocinative and Inductive*, volume 1 (London: John W. Parker).
Miller, G. (2000) *The Mating Mind: How Sexual Choice Shaped the Evolution of Human Nature* (New York: Doubleday).
Nagel, T. (1986) *The View from Nowhere* (Oxford: Oxford University Press).
Neander, K. (1995) ‘Pruning the Tree of Life’, *British Journal for the Philosophy of Science*, 46: 59–80.
Nozick, R. (1993) *The Nature of Rationality* (Princeton: Princeton University Press).
——— (2001) *Invariances: The Structure of the Objective World* (Cambridge, Mass.: Harvard University Press).
Papineau, D. (2000) ‘The Evolution of Knowledge’, reprinted in his *The Roots of Reason* (Oxford: Oxford University Press, 2003), 39–82.
Peacocke, C. (1992) *A Study of Concepts* (Cambridge, Mass.: The MIT Press).
Plantinga, A. (1993) *Warrant and Proper Function*. (Oxford: Oxford University Press).
Rips, L. (1994) *The Psychology of Proof: Deductive Reasoning in Human Thinking* (Cambridge, Mass.: The MIT Press).
Ryle, G. (1950) ‘”If,” “So,” and “Because”,’ in M. Black (ed.), *Philosophical Analysis* (Ithaca: Cornell University Press), 323–340.
Schechter, J. (2010) ‘The Reliability Challenge and the Epistemology of Logic’, *Philosophical Perspectives*, 24, Epistemology, 437–464.
——— and D. Enoch (2006) ‘Meaning and Justification: The Case of Modus Ponens’, *Noûs*, 40: 687–715.
Skyrms, B. (1999) ‘Evolution of Inference’, in T. Kohler and G. Gumerman (eds.), *Dynamics in Human and Primate Societies* (Oxford: Oxford University Press), 77–88.
Sloman, S. (1996) ‘The Empirical Case for Two Systems of Reasoning’, *Psychological Bulletin*, 119: 3–22.
Sober, E. (1981) ‘The Evolution of Rationality’, *Synthese*, 46: 95–120.
——— (1984) *The Nature of Selection: Evolutionary Theory in Philosophical Focus* (Cambridge, Mass.: The MIT Press).
Stich, S. (1990) *The Fragmentation of Reason* (Cambridge, Mass.: The MIT Press).
Sterelny, K. (2003) *Thought in a Hostile World: The Evolution of Human Cognition* (Oxford: Blackwell Publishers).
——— (2007) ‘Social Intelligence, Human Intelligence and Niche Construction’, *Proceedings of the Royal Society, Series B*, 362: 719–730.

Stroud, B. (1981) ‘Evolution and the Necessities of Thought’, Reprinted in his *Meaning, Understanding, and Practice: Philosophical Essays* (Oxford: Oxford University Press, 2000), 52–66.

Tennant, N. (1987) *Anti-Realism and Logic* (Oxford: Clarendon Press).

Williamson, T. (2003) ‘Understanding and inference’, *Proceedings of the Aristotelian Society, Supplementary Volume*, 77: 249–293.

——— (2007) ‘Philosophical Knowledge and Knowledge of Counterfactuals’, *Grazer Philosophische Studien*, 74: 89–123.