

CHANGING OUR LOGIC: A QUINEAN PERSPECTIVE

Mind (2023). doi.org/10.1093/mind/fzad010

Michael Devitt and Jillian Rose Roberts¹

The Graduate Center, The City University of New York and Princeton University

1. Introduction

Can we change our logic and if so how? In ‘The Question of Logic’ (this volume), Saul Kripke takes a certain message about this from Lewis Carroll’s famous paper (Carroll 1895), ‘What the Tortoise Said to Achilles’. That message concerns certain basic principles of valid inference, such as those of universal instantiation (UI) and *modus ponens* (MP).² Romina Padró puts the point this way in ‘The Adoption Problem and the Epistemology of Logic’ (this volume):

Adoption Problem (AP): certain basic logical principles cannot be *adopted* because, if a subject already infers in accordance with them, no *adoption* is needed, and if the subject does *not* infer in accordance with them, no *adoption* is possible. (##3)³

Padró introduces a character, Harry, who has never made an inference in accordance with the UI pattern but has been brought to accept the UI principle. Padró elaborates helpfully on the problem that she and Kripke have in mind by describing what it would be for Harry to ‘adopt’ UI:

By ‘adopt’ here we mean that he picks up a way of inferring according to UI...*on the basis* of the *acceptance* of the corresponding logical principle... ‘adoption’ is a two-phase process: *Phase 1* is given by Harry’s *acceptance* of the UI principle, and *Phase 2* consists in the development, *in virtue of Phase 1*, of a practice of inferring in accordance with the UI principle. (##3-4)

Padró takes the case of Harry to be a *reductio* of the idea that one could adopt an inferential practice by accepting a basic logical principle.

¹ The first version of this paper, under the title ‘The Adoption Problem: A Quinean Solution’, was delivered by Devitt at a conference, ‘The Nature of Logic’, at University of York (UK), January 2016, in response to papers by Saul Kripke and Romina Padró. The three of them more or less repeated this performance at the Graduate Center of CUNY, April 2019, observed by Roberts. This prompted her to take an ‘independent study’ course with Devitt on the Adoption Problem in Fall 2019. Roberts convinced Devitt that his naturalistic ‘solution’—roughly what is entertained in §6—does not work. They decided to work together to find ones that would work, largely based on ideas from Roberts, and write a joint paper. Those developments are in §7 and §8.

² The UI principle may be roughly stated this way: that all *F*s are *G* implies any particular *F* is *G*. For example, that all humans are mortal (and Socrates is a human) implies Socrates is mortal. The MP principle may be roughly stated this way: if *P* then *Q*, and *P*, together imply *Q*. For example, if Socrates is a human then he is mortal, and Socrates is a human, together imply he is mortal.

³ All otherwise unidentified citations of Kripke and Padró are to their papers in this volume.

Kripke and Padró are particularly critical of what they see as Quine's view of the matter. We shall present and defend a picture of AP, indeed of the broader problem of 'changing our logic', that is clearly inspired by Quine's naturalistic stance. We make no claim that the picture is Quine's.

AP is built around the important distinction between accepting a principle and following a practice, a *skill*. We base our naturalistic solution on what science tells us about this distinction. Now this solution to AP, *as stated*, can obviously be resisted by saying that it is excluded by some presumed background rules for a solution. Then the issue becomes one about the appropriateness of these presumed rules. We think that there is no basis for excluding our solution. Still, there remains a respect in which AP really does seem to be insoluble. This does strike us as a bit surprising, but we wonder about its interest given the respect in which AP is soluble.

In the next section, we define the two questions raised by AP. The Quinean answers that we shall present utilize psychological views of what it is for a process to be rule-governed, discussed in §3, and of how we learn to be governed by new rules, discussed in §4. We turn to the Quinean picture in §5 to §9.

2. The Acceptance Question and the Adoption Question

The important distinction that underlies AP is as follows. On the one hand, there are *the rules* that govern *the processes* by which a person forms beliefs, the rules that constitute what Hartry Field calls her 'evidential system' (Field 1996; Field 1998). These must include rules for forming beliefs from perceptual experiences and the logical rules that concern us here, rules for inferring one belief from another. On the other hand, there are *theories about such rules*. Thus, it is one thing for a rule of inference *R*—for example, UI—to be among the rules that govern a person's reasoning, it is another thing to theorize that *R* does so govern.⁴ And it is another thing again to theorize about whether our inferences *should be governed* by *R*. Is *R* a *good* rule in that it gives epistemic warrant to its conclusion when it operates on true premises? Should it be replaced by another rule *R'* that gives a better warrant? In sum, there is the *epistemic practice* of a reasoning that is governed by *R*, and this contrasts with two types of *epistemological theorizing* about *R*. The first type of theorizing is a piece of *descriptive* epistemology, the second, a piece of *normative* epistemology.

In light of this, we distinguish two questions for changing a logic. One of them concerns a person *coming to accept* a new basic logical principle; this acceptance comprises understanding the principle, believing it is true, and thereby accepting that inferential practices should be governed by a new rule *R'*, perhaps replacing an old rule. The other concerns a person 'adopting' *R'* in the sense of *coming to be governed by R'* in her inferential practices. Now it is obvious that, technical difficulties aside, these changes could be brought about by mere surgery, bumps on the head, and so on. That is boring. The questions that concern us are about *rational* changes. Can she rationally come to accept the new basic principle? Can she, 'on the basis of', 'in virtue of', accepting a new basic principle, rationally change her

⁴ In actual fact, we surely do not infer simply in accord with UI, as Gilbert Harman has made plain (Harman 1999, pp. 18-23): if we believe that all *F*s are *G* and that *a* is an *F* then we *might* indeed infer that *a* is *G* but we might be so convinced that *a* is not *G* that we abandon our belief that all *F*s are *G*. The relations between psychological processes of inference, even of good inference, and logical implications are complex. Still, UI is surely involved in some of our inferences and those inferences are good only because UI is valid.

practices so that her inferences are governed by R' , as the principle recommends? So we have the following two questions, each with two parts:

Acceptance Question: (i) Can a person *rationally* accept a new basic logical principle, accepting that inferences should be governed by a new rule R' (perhaps replacing an old rule R)? (ii) If so, how?

Adoption Question: (i) Can a person, on the basis of accepting a new basic logical principle, *rationally* change her practices so that her inferences are governed by R' (perhaps replacing an old rule R)? (ii) If so, how?

The Kripke-Padró AP is the claim that the answer to the Adoption Question (i) is 'No'. We shall answer 'Yes'.

Padró follows her characterization of adoption (§1) with the surprising claim that two possible solutions are 'automatically excluded' by it: 'mere causal connections' and the use of 'training' (##4). We disagree. In particular, the characterization does not exclude, let alone 'automatically' exclude, training. This is important because our Quinean picture proposes such a solution. We shall discuss Padró's exclusions in §4. First, we turn to psychology for a scientific view of what it is for our inferences to be, and rationally become, rule-governed.

3. Rule governing and skills

The practice of inferring is a cognitive skill, indeed, the pre-eminent cognitive skill. It is the skill of moving in a (hopefully) rational way from one belief to another. When we talk of this practice being rule-governed we have in mind the way skills are rule-governed. What way is that? To answer this question, we should look to the psychology of skills. The consensus view in psychology is that skills are pieces of *procedural knowledge*, which can be *roughly* identified with what the folk call 'knowledge how'. Procedural knowledge is sharply distinguished from *declarative knowledge* which can be identified with what the folk call 'knowledge that'. This distinction is fundamental to cognitive psychology (Anderson 1980, p. 223).⁵

What is declarative knowledge? The key thing to note for our purposes is that there is a consensus in psychology that it involves a *conscious representation* of what is known. Thus, psychologists think that a subject has declarative knowledge of the processing rules for a task only if she consciously represents the rules. So the person who has declarative knowledge that R is a rule of arithmetic must represent that fact *in her 'central processor'*, the home of her propositional attitudes. Hence the identification of declarative knowledge with the folk's knowledge-that.

Before considering the nature of procedural knowledge, it is helpful to mention another distinction.⁶ It comes largely from computer science. It is the distinction between processing rules that govern by being represented and applied and those that govern by being simply embodied without being represented. This is a distinction between two ways in which certain processing rules might be real in an object, two ways in which the rules might be embodied in it. Neither of these ways should be confused

⁵ (Tenison and Anderson, 2015) is a recent example of the distinction in use.

⁶ See (Devitt 2006, pp. 45-52) for a detailed discussion of this distinction.

with a situation where an object simply behaves *as if* it is governed by those processing rules. For that situation is compatible with those rules *not* being embodied in the object at all.

A simple old-fashioned mechanical calculator provides a nice example of something governed by rules that are embodied without being represented. When the calculator adds it goes through a mechanical process that is governed by the rules of an algorithm for addition. But the rules are ‘hardwired’ not represented in the calculator. In contrast, the operations of a contemporary general-purpose computer are partly governed by rules of a program that are represented in its RAM and applied. Yet those rules can govern the operations of the computer only because there are other rules that are unrepresented but built into its hardware that enable the represented rules to govern. And, note an important generalization: any processing rule that governs the behavior of one object by being represented and applied could govern that of another by being embodied without being represented.

What is procedural knowledge and hence what are skills? As already noted, psychologists agree that this knowledge is *not* declarative but, beyond that, we have a long way to go in answering this question (Schacter 1999, p. 395; Sun 2003, p. 698; Burgin 2016, p. 49). The psychological literature reveals a range of interesting ideas but no rational basis at this time for a sweeping acceptance or rejection of the ideas of one or other theoretical camp. Do processing rules involved in procedural knowledge and skills govern by being represented and applied? Perhaps so, but we think that the weight of evidence counts against this; rather, the rules govern by being embodied without being represented (Devitt 2006, pp. 212-216). However, even if the rules are represented, they are not represented in propositional attitudes but in some sub-central module of the mind. As Mark Burgin points out in his encyclopedic study of structures and processes of knowledge, the procedural knowledge that constitutes a skill is stored in a different part of the brain from declarative knowledge (Burgin 2016, p. 51).

There is much that psychology has not yet discovered about skills, but what it has discovered should discourage over-intellectualizing our inferential skills. So, we doubt some of the views Padró assumes for her *reductio*, many of which she believes are widely accepted in epistemology:

Competing justification proposals have generally assumed that...[a] (tacit) grasp or acceptance of [logical principles] is supposed to guide our inferential practices. It is because a thinker has accepted or grasped the UI principle itself, for example, that she is in position to perform transitions that accord with UI. (##2)

She elaborates on this dependence of our inferential practices on our acceptance of logical principles, insisting that ‘the principle’s propositional content’ must not be ‘superfluous to the inferential act’ (##7). There is ‘a constitutive relation between the principle and the inferential pattern’ (##7). As noted, it is possible, though we think unlikely, that the principles are sub-centrally represented and applied in inferences. So, one might say, though we would not, that a ‘(tacit) grasp’ of principles, *in that sense, may* guide our inferential practices.⁷ But that would not make the principles and their contents *constitutive* of an inferential practice. Rather that (tacit) grasp would be the *contingent* way that the practices are

⁷ Linguistics has lent encouragement to such talk. Thus, linguists often talk of competent speakers of a language standing in an unconscious or tacit propositional attitude to the rules of the language which are represented in her language faculty. We think that there is no significant evidence for this view and, given what else we know, it is implausible (Devitt 2006: 87-121, 195-272).

causally implemented in humans; implementation might well be different in chimps, dolphins, and crows. Finally, Padró thinks that her tacit-grasp view ‘has considerable intuitive appeal’ because it is not clear how else we would explain our capacity ‘to produce a potentially infinite number of UI inferences we never considered before’ (##8). But we would have a way: embodied but unrepresented rules could explain this ability. And they likely do.⁸

Two further points. (1) None of this denies that inferential acts are, as Padró rightly insists, ‘voluntary and under our rational control’ (##8). Reasoning, like touch typing and bicycling, is an intentional activity. But what we intentionally set in motion is simply a process governed by sub-centrally embodied rules (whether represented or not). (2) The premises that feature in a person’s rule-governed inference are reasons for her conclusion *simply in virtue of so featuring*. There is no requirement that her reasoning be accompanied by the *further thought* that ‘the conclusion [is] supported by the premise(s)’, contrary to the tacit-grasp view that Padró finds so appealing.

4. Learning a skill

Rules that govern the operation of a computer by being represented and applied are in the ‘software’ and can be fairly easily changed: just change the program. Rules that govern by being simply embodied are in the ‘hardware’ and it takes a bit of engineering to change them: for example, replace the Intel chip. There is an analogous story for humans. Rules that govern by being represented in our central processor and applied – for example, the diplomatic rules for a state dinner – are in the software and can be changed easily by changing beliefs. In contrast, the rules that govern a skill are embodied in the ‘wetware’ and are typically quite hard to change. Still, such changes in the wetware are common: they occur every time we acquire a skill; for example, when we learn to change gear, ride a bicycle, play chess, or, let’s suppose, reason better.⁹

How *do* we change the wetware to acquire a new skill? Psychology’s fundamental distinction between declarative and procedural knowledge is accompanied by another distinction in attempting to answer this question. The distinction is between *explicit learning* and *implicit learning*. Explicit learning is a ‘top-down’ process that starts from declarative knowledge. Consider, for example, explicitly learning the skill of high jumping using ‘the Fosbury Flop’.¹⁰ We start with declarative knowledge of how to jump in this way, yielding a series of instructions: ‘Approach the bar at your greatest controlled speed, plant your foot at approximately 20 degrees to the bar at take-off, thrusting your arms and knee as hard as possible upward, while keeping your eye on the bar and arching your back; at the last second, pull your back leg over the bar, and land on the base of your shoulder’. We follow the instructions and, *with practice*, what the psychologists call ‘proceduralization’ occurs and all these processes become automatic. The role of the declarative knowledge in this process is as a starting point *to train* in the skill. However, once we have acquired the skill, the declarative knowledge may well be lost. Thus, Caitlan

⁸ Padró claims optimistically that her tacit-grasp view is so appealing that ‘it has hardly been questioned’ (##8-9). Well, we question it and think many psychologists do too, implicitly at least.

⁹ We should, of course, not allow this metaphorical talk of ‘software’ and ‘wetware’ to distract us from the obvious truth that both embodiments are in the brain.

¹⁰ This method of high-jumping is named after Dick Fosbury who used it first to win the gold medal at the 1968 Olympics. Since then it has been used by all leading high jumpers.

Tenison and John Anderson (Tenison and Anderson 2015) describe the final stage of acquiring the arithmetical skill exemplified by multiplying 3 by 5 as follows:

With enough practice the learner can acquire a production that simply produces the answer 15 in response to the problem without querying declarative memory. While strategy choice occurs in the two prior phases, during the final autonomous phase, time is not spent choosing a strategy; rather, answering the question effectively become a reflex. (Tenison and Anderson 2015, p. 3)

In contrast, implicit learning is a ‘bottom-up’ process: we observe, practice, and ‘just pick the skill up’; think of learning to ride a bicycle or throw a Frisbee. There is much evidence that a lot of skill learning is implicit; see, for example, the evidence cited by Sun *et al* (Sun *et al* 2001) that ‘individuals may learn complex skills without first obtaining a large amount of explicit declarative knowledge... and without being able to verbalize the rules they use’ (Sun *et al* 2001, p. 207).

We earlier defined the adoption that concerns us as a matter of rationally coming to be governed by inferential rule *R'* on the basis of accepting a new basic logical principle. The accepted logical principle is a piece of declarative knowledge. So if this skill of reasoning by *R'* can be acquired in any interesting way on the basis of accepting a logical principle it must be acquired by *explicit* learning. In explicitly learning a reasoning skill, a person has declarative knowledge about *R'* and she uses this as a starting point to *train* to reason by *R'*, just as a person uses declarative knowledge about the Fosbury Flop to *train* to Flop. *Any positive and interesting answer to the Adoption Question must appeal to explicit learning.*

In explicitly learning a skill, what are the mechanisms that get us from declarative knowledge to training and on to procedural knowledge? That’s another very difficult question that psychologists struggle with (Tenison and Anderson 2015).

Whatever the process of explicitly learning a skill, the process is *simply* a causal one (which is not to say that it is a *simple* causal one!). So the picture for logic is that a person’s declarative knowledge that *R'* is a good inference rule *causes* her over the period of training to infer by *R'*. No *other* ‘basis’ is required for rule adoption (*pace* Padró’s view stated below). And the process is *rational*. Whenever we come up with a theory of a better way to perform some activity, whether to high jump using the Fosbury Flop or to infer using *R'*, it is rational for a person engaged in that activity to train to perform that way.

It is time to consider Padró’s ‘exclusions’. Refuting her AP requires us to show how someone like Harry could acquire the cognitive skill of inferring by UI ‘on the basis of’, ‘in virtue of’, the acceptance of the UI principle. Yet, having stated AP, Padró immediately claims that some possible solutions are ‘automatically excluded’. We have already noted that this is not so (§2): these exclusions are additional. Furthermore, they are mistaken. Here are the exclusions:

First, mere causal connections, where the relation between Harry’s acceptance of the principle and his inferring in accordance with it could be fortuitous, do not count as a case of adoption. His acceptance of the principle should *guide* Harry’s UI inferential practice. Second, the possibility of developing the UI practice by means of different kinds of training or psychological conditionings – even when accompanied by the acceptance of the UI principle – is excluded. Harry’s acceptance

of the principle should be his *sole* reason and explanation for inferring in accordance with it.
(##4)

(a) Connections between a skill and its ‘principles’, if any, will indeed be ‘mere causal’ ones (§3). (b) There is nothing ‘fortuitous’ about learning a skill by intentionally training to acquire it. (c) The only way to get from *knowledge about* a skill to *having* the skill is by training; that’s the route from *knowledge-that* to *knowledge-how*. (d) Most important of all, *Padró’s exclusions, would make all explicit skill learning impossible*. Not only could we not learn to infer ‘on the basis of’, ‘in virtue of’, accepting a principle (declarative knowledge), we would not be able to learn any skill on that basis: not to Fosbury Flop, touch type, assemble Ikea furniture, play chess, etc. For, *accepting a principle is not the ‘sole’ explanation of acquiring any skill*. We know this from science (and commonsense). In sum, *Padró’s exclusions make AP uninteresting. The exclusions should be rejected.*¹¹

None of this is to say that *R’* could not be acquired by *implicit* learning. Indeed, it seems quite likely that as a child grows up she does acquire some reasoning skills by implicit learning, by ‘experiencing what works’. But, of course, this would not be acquisition as a result of accepting a logical principle, which is the concern of AP.

Utilizing these views of what it is for a process to be rule-governed and of how we learn to be governed by new rules, we turn now to our naturalistic Quinean answers first to the Acceptance Question and then to the Adoption Question.

5. A Quinean answer to the Acceptance Question

We start with some relatively uncontroversial background. Millions of years ago, before anyone had any logical principles, our hominid ancestors came to embody logical rules in their evidential systems, rules that governed the inferences they used to build their ‘webs of belief’ in the face of experience. There is good reason to think that the same is true of other animals, particularly primates and dolphins.¹² These embodied rules evolved over time, probably by a series of adaptations. Despite the common inheritance

¹¹ In ‘Kripke, Quine, the “Adoption Problem” and the Empirical Revision of Logic’ (this volume), Paul Boghossian and Crispin Wright raise the thought that the very understanding of ‘all’ requires that one infer by UI rule. We rather doubt this thought, but suppose it were so. Then Harry, who does not infer by the UI rule, would be unable even to understand a statement of the UI principle, because it uses ‘all’ or an equivalent. So, he could not *accept* the UI principle in the relevant way: the acceptance requires the understanding. So, obviously, he could not adopt the UI rule *on the basis of* accepting the UI principle. But suppose that Harry uses ‘all’ just as we do *apart from* failing to employ it in UI inferences. He still infers from a belief that *each* item of a group has a certain property to a belief that *all* items of the group have that property. Since he does not instantiate, we might say he has an ‘instantiation-neutral’ (Boghossian and Wright, this volume) understanding of ‘all’. When someone who uses ‘all’ in the standard way instructs Harry about the UI principle, Harry comes to accept a weaker principle that he would express using ‘all’ in this instantiation-neutral sense. So our Adoption Question still remains to be answered.

¹² The nature of these beliefs is an interesting issue, of course; for example, do they involve ‘languages of thought’? We take no stand on this very difficult issue.

of embodied rules among members of the human species we expect genetic variation in the rules used by the members as in everything else. And probably what came to be embodied in each adult member was not solely a matter of an inherited nature but partly a matter of nurture: the embodied rules were partly the result of implicit learning, just as the embodied the rules for many a skill are.

There was no interesting sense in which our ancestors ‘chose’ their governing rules any more than they chose any other of the properties the human lineage evolved to have. In brief, our ancestors got their rules largely by evolution and implicit learning not by a ‘rational’ process.

Turn now to part (i) of the Acceptance Question. Only very recently in the human lineage – starting we are told with the ancient Greeks about 3000 years ago – have a very few humans, led by logicians, ‘chosen’ to make logical principles part of their webs of belief. Before that, for millions of years, hominids have embodied rules that govern their inferences without any theory about which of those rules are good. Thus, suppose that UI is one such embodied rule. Then, on this picture, it is a brute fact about our species that we infer from universal beliefs to their instances. We were doing this long before anyone noticed that we were doing it, came to expect us to do it, or discovered the principle that we ought to be doing it.

The acceptance of logical theories did not, of course, stop with the ancient Greeks. Aristotelian logic was followed by modern logic; classical logic has been confronted with alternative logics. We have come up with modal logics, tense logics, and so on. And we have come up with theories for non-deductive ‘ampliative’ reasoning: probability theory, non-monotonic logics, and so on. So, *many* logical theories have been accepted and some have been replaced. We suppose that just about everyone agrees that this process of acceptance and replacement was a rational one, even if sometimes mistaken. So we have a swift answer to part (i) of our question: a person certainly *can* come rationally to accept a new basic logical principle because we *have*. And UI is a principle humans accepted quite recently.

Next, consider part (ii) of the Acceptance Question. *How* do we do come to accept basic logical principles? *What sort* of rational process is this? This is, of course, a notoriously difficult and controversial question. The philosophically most popular answer is that we accept these principles by some non-empirical process of rational insight or a priori justification. To the Quinean, this is not a reputable answer because we do not have faintest idea about the nature of this non-empirical process of knowing. We are told what it is *not*—it is *not empirical*—but we have never been given anything close to a satisfactory account of what it *is*. It is left quite mysterious (Devitt 2011).

If the acceptance process is not a priori then it must be empirical. But how can it be? Suppose that the new theory claims that R' is a better rule than R and so should replace it in our logic. Here is a first stab at showing how this could be shown empirically to be rational. We observe that inferring according to R' is more successful than inferring according to R ; it is more successful in that it more reliably takes us from empirical truth to empirical truth.

Objection. How do we tell that R' is more successful than R ? To tell this we must of course go through a process governed by our present evidential system S . How else? I can see, perhaps, how this process could be rational if neither R' nor R is already in S . But suppose that R' is in S . Then we seem to have ‘rule circularity’: R' is used to justify its own success and hence acceptance. Suppose, on the other

hand, that R is in S but R' is not. Then we seem to have ‘self-defeat’: R is used to justify its failure and hence overthrow. (And what about if *both* R' and R are in S , an ‘inconsistent’ system?!)

Response. The first move in a response appeals to Quine’s famous metaphor of Neurath’s boat. Quine likens our web of belief to a boat that we continually rebuild whilst staying afloat on it. We can rebuild any part of the boat—by replacement or addition—but in so doing we must take a stand on the rest of the boat for the moment. So we cannot rebuild it all at once. Similarly, we can justify or revise any part of our knowledge but in so doing we must accept the rest for the time being. So we cannot justify or revise it all at once. So the claim that one of S ’s rules is successful or not could be supported by an argument that uses *other* rules of S but not that rule itself or its proposed replacement.

We are not so optimistic as to suppose that this Neurathian move is always available, particularly when we consider the basic logical principles (let alone the rules within our evidential system that govern its own change). So we do think that rule-circularity is a worry but it is not clear just how big a worry it is (Papineau 1993; Psillos 1999; Boghossian 2000; Boghossian 2001). And a comfort for the naturalist is that it is just as much of a worry for the apriorist; or so it has been argued (Devitt 2011, pp. 281-282).

Here is a reason not to worry about a potential problem of self-defeat. In §6 below we present examples of theory acceptance leading to rule adoption. Now each such theory acceptance was, of course, brought about by the exercise of an evidential system. So,

that system was used to establish an epistemological thesis that led to the system’s replacement. These examples give us good reason to think that an evidential system could be used rationally to undermine itself. Accepting the non-epistemological part of our web and governed by S as usual, we find [the theory that S is good] wanting and so replace it and the system S that it recommends. (Devitt 2005, p. 110)

We might say that this is a kind of *reductio* of the theory that S is good. But there is still a reason to worry, as Paul Boghossian and Crispin Wright point out in ‘Kripke, Quine, the “Adoption Problem” and the Empirical Revision of Logic’ (this volume). A decision that S is not good is *rational* only if it is arrived at by sound reasoning. So our decision that S is not good, *which was reached using* S , undercuts the rationality of that very decision.

There might be a way out.¹³ Suppose that our decision that S is not good arises from our view of its rule R : R needs to be replaced by R' . Then we would have a way out if our decision that S is not good is supported by our new system using R' instead of R . So, whether we use R or R' , either way we come to the conclusion that R' is better than R . However, if using R' does not support the conclusion, then it seems that we must rest on the *reductio*.

How far can our logical theory go in replacing rules? Could R be a basic rule like UI or MP? In Quinean ‘anti-exceptionalism’ about logic, the goodness of any rule, even UI, is in principle open to challenge. But this is not to say that we should seriously consider changing it; it is not to say even that such change is ‘conceivable’. And to say that, when confronted with what appears to be a white raven, we have ‘the choice’ of abandoning UI is not to say that abandoning UI in these circumstances has any

¹³ Thanks to Daniel Boyd.

merit. It is just to say that our theories of good reasoning, like our theories of anything else, must answer to the tribunal of experience.

Finally, we respond to the following passage from Kripke:

Quine...says that the logical laws are ‘simply certain further statements of the system, certain further elements of the field’, just like any others. If one wishes to modify one’s system one can choose to modify either the particular hypotheses like ‘all ravens are black’ or the logical laws. The impression you will get from that kind of discussion is that someone would not be intellectually committed to asserting that ‘this raven is black’ just because all ravens are black. But he is committed to ‘this raven is black’ if, *in addition*, he accepts various logical laws, in particular: All universal statements imply their instances....If we understand [the law of universal instantiation] as a hypothesis or belief that we ‘adopt’ in order to (indirectly) determine its empirical impact, we should conclude that its scientific fertility is *zero*. It will not lead to a single prediction. (##13)

Perhaps there can be a Kumbaya moment here. For, we think that Quineans should acknowledge that *logical principles are indeed different* from other propositions and should agree with Kripke that ‘logic...cannot be just like geometry’ (##17). And the ways logical principles differ go against the above impression about commitment. Those differences are (1), in the way logical principles affect other propositions and (2), in the way other propositions, and experience, affect them. Here’s a stab at a Quinean picture of these differences.

Concerning (1), let P be some non-logical, uncontroversially empirical, proposition in a person’s web of belief, say, ‘all ravens are black’. P can *bear directly on other such propositions* that constitute the web by the person’s exercise of her evidential system S : the system will treat P as inconsistent with some propositions, supporting others and so on. So, assuming S includes the UI rule, P does commit her to ‘this raven is black’. And she has this commitment even if she has never thought of the UI principle. The story for a logical principle L , say, the UI principle, is very different from this one for P . The person’s belief in the UI principle cannot bear *directly* on her non-logical beliefs:¹⁴ Regarding such non-logical beliefs, the UI principle will indeed ‘not lead to a single prediction’. It can bear on those beliefs only by causing her, in ways to be discussed, to infer according to the UI rule. That causal process can change S , the epistemic system that relates her non-logical beliefs to one another. Hence the UI principle can bear *indirectly* on her non-logical beliefs when the system is exercised. That is the way it can be scientifically fertile.

Concerning (2), exercise of the person’s evidential system S brings experience to bear on P via its relations to many other beliefs—the Duhem-Quine thesis—but still, in a sense, *directly* confirming or disconfirming P . This is not the case with a logical principle like UI. When the person’s exercise of S in the face of recalcitrant experience leads to propositions that are in conflict with each other, she can *contemplate* changing S itself. So she might contemplate not inferring by UI when experiencing what seems to be a white raven, thus avoiding the unwelcome conclusion that all ravens are not black (though

¹⁴ Other than ones like ‘Jill says that L is false’ which refer to L itself. An evidential system will treat such beliefs as in tension with or supporting L . In this way, L bears more directly on a belief that refers to L than it does on a belief that does not.

she surely wouldn't and shouldn't). To contemplate this is to contemplate rejecting the UI principle, the principle that recommends the UI rule.¹⁵ Thus recalcitrant experience can be brought to bear on the UI principle *indirectly*. This amounts, in Quine's words, to 'amending certain statements of the kind called logical laws' (Quine 1951, p. 40). Quine asks 'what difference is there in principle between such a shift and the shift whereby Kepler superseded Ptolemy, or Einstein Newton, or Darwin Aristotle?' (Quine 1951, p. 40). Well, the indirectness of experience's impact on a logical law is certainly a difference but perhaps it does not count as a difference 'in principle'.

We turn now to the Adoption Question.

6. A Quinean self-training answer to the Adoption Question

Consider part (i) of the Adoption Question. Can a person, on the basis of accepting that R' is a good basic rule of inference, rationally change her practices so that her inferences are governed by R' ? Now, as we have noted, it is indubitable that we have adopted new logical theories. *But have the new theories changed our logical practices?* It may be difficult to tell because we know so little about S , so little about the evidential system that actually governs our reasoning; descriptive epistemology is in its early days. Still, our empirical hypothesis is that our theoretical changes have caused lots of changes in our practices.

When it comes to non-deductive ampliative reasoning this hypothesis does not seem bold. First, it seems undeniable that over the last three centuries we have developed better and better theories of how sciences should go about finding the truth. We have learnt a vast amount not only about the world but also about how to learn about the world. Much of the education of the young scientist is in these methodologies: think of physics and psychology, for example. And we take it as rather obvious that the improved methodologies that have been taught have been then *used*: those young scientists practice what they have been preached. A good deal of the impressive scientific progress over recent centuries has been the result of this application of improved methodologies. Changes in methodological theories have led to changes in our evidential systems.

Second, our theories tell us that certain sorts of ampliative inferences are bad: counter-induction; ignoring the base rate in probability calculations, and so on. Yet psychologists tell us that just about everyone is prone to make these mistakes. Those in the know about all this try, and sometimes succeed, in avoiding these mistakes. Once again we have theory change leading to changes in our reasoning practices.

But what about deductive logic? Our hypothesis that even our deductive practices have changed as a result of logical theories may seem rather bold. Still, it gets a lot of support from anecdotal evidence. People who have learnt logic commonly claim that doing so improved their reasoning. Philosophy departments include claims to this effect as part of their pitch for students. Still, arguably, all these improvements in our deductive reasoning are matters of dropping bad rules—for example, the fallacy of affirming the consequent—or 'reinforcing' good rules. The reinforcement we have in mind is as follows. It is likely true for (almost?) any of the logical principles that everyone's reasoning is sometimes governed by the rules the principle specifies it should be governed by and sometimes not: our skill at reasoning is

¹⁵ The contemplation of this shows that S must include not only processes for moving from one proposition in the web to another but also processes for assessing the goodness of inferences.

imperfect just as our skill at anything else is. Learning logic likely causes us to be governed by these good rules more often. But this type of improvement in our reasoning is not the same as coming, as a result of adopting a basic principle, to be governed by a reasoning rule that *we were not governed by at all*. And it is the latter type of improvement that Padró thinks impossible. This is what Harry's problem is supposed to demonstrate (##4-6).

We would like Harry to perform the following UI inference:

Argument A

- (A1) All the animals in the movie talk.
- (A2) Alex is an animal in the movie.
- (A3) So, Alex talks. [from (A1) and (A2) by UI]

But Harry doesn't infer by the UI rule and so does not reach (A3). To help Harry out we tell him about the UI *principle*. Harry believes us, thus accepting the principle: 'All universal statements imply each of their instances.' But how could his accepting it help Harry with his inference? How could it result in his adopting the UI rule and hence making the desired inference? For it to help, Padró presumes, he would have to begin by reasoning as follows:

Argument B

- (B1) All universal statements imply each of their instances
- (B2) 'All the animals in the movie talk' is a universal statement
- (B3) So, 'All the animals in the movie talk' implies each of its instances [from (B1) and (B2) by UI¹⁶]

Once Harry has reasoned his way through Argument B, he can use (B3) to begin the desired inference, Argument A. But, as Padró points out, Harry will never get through Argument B unless he *already* infers by UI. For, the inference from (B1) and (B2) to (B3) is a UI inference. So, she concludes, accepting the UI principle cannot result in adopting the UI rule. Similar arguments can be constructed with other basic logical principles (such as MP). So, says Padró, 'Harry illustrates our rough approximation to the AP: certain basic logical principles cannot be *adopted* because, if we do not *already* infer in accordance with them, accepting them would lead us nowhere' (##6). That is her *reductio*.

This brings us to part (ii) of the Adoption Question. If we look to the account of explicit skill learning in §4, we see that this failed attempt to get Harry to adopt the UI rule presumes the wrong view

¹⁶ As Padró notes, though we refer to this form of inference as UI for simplicity's sake, in formal first-order logic, a logician should formalize it as universal *modus ponens* (##5), $(\forall x)(A(x) \rightarrow T(x)), A(a) \therefore T(a)$. Moreover, it is often formalized by semanticists without an embedded conditional, as $(\forall x: A(x)) A(a) \therefore T(a)$, and so apparently not requiring *modus ponens*. How it is formalized, however, does not change what Harry *has to do* in order to *correctly* apply the general principle, (B1). Harry will need to apply (B1) in a conditional form: that *if* there is a universal statement of interest, *then* he should conclude that the universal statement implies each of its instances. This point is covered more thoroughly in §7.

of how accepting the UI principle *might* result in adopting the UI rule. The logical principle is declarative knowledge, whereas being governed by the rule is procedural knowledge. Procedural knowledge is not declarative knowledge and so adopting the principle, indeed adopting any declarative knowledge, could not alone *ever* be sufficient for the procedural knowledge of reasoning. So, how does declarative knowledge of a skill help acquire the procedural knowledge of the skill? The declarative knowledge about a skill can result in the procedural knowledge *only by yielding the instructions used in explicit learning to train someone in the skill*. A person explicitly learning the skill must follow these instructions *repeatedly* to train in and acquire the skill.

Return to our earlier example. Declarative knowledge about the Fosbury Flop yields a series of instructions on what to do to high jump: ‘Approach the bar at your greatest controlled speed, plant your foot at approximately 20 degrees to the bar at take-off’ and so on. A person wanting to Flop follows these instructions and jumps. *After many repetitions* of this behavior, the person acquires the skill of Flopping. If Harry is to adopt UI, he must go through such an explicit-skill-learning process. Declarative knowledge of UI yields an instruction on what to do when you want information about a certain *F* entity, for example, about the animal Alex in the movie. The instruction is along the following lines: ‘infer an instance about that entity from a universal you believe about *F*s’, e.g. about Alex. Harry believes the universal, ‘all the animals in the movie talk’, follows the instruction, and so infers the instance, ‘Alex talks’. *After many repetitions* of this sort of behavior, Harry acquires the skill of UI reasoning. Following the Fosbury Flop instructions requires quite a few skills: for example, recognizing the bar and knowing how to arch your back. Similarly, following the UI instruction requires recognizing a universal statement and knowing what an instance of it is. But none of these required skills seems to be a reasoning skill, whether at UI or anything else.

7. The failure of the Quinean self-training answer

So, with knowledge of the UI principle, Harry could rationally put himself through this training process to acquire the cognitive skill of inferring by the UI rule. We seem to have a neat Quinean solution to AP. But, sadly, we do not. The give away is in what precedes the instructions. Thus, in the Fosbury Flop case, we say, ‘when you want to high jump’, before giving the instructions. But the instructions themselves *do not specify the appropriate circumstances for attempting to Flop*: you should be at an athletic field and want to high jump. Obviously, only when you are in those circumstances should you follow the instructions; you don’t try to approach a high jump bar when you are in your office or want to open a beer. Similarly, in the UI case, we say, ‘when you want information about a certain *F* entity’, before giving the instructions. The instructions themselves *do not specify these circumstances for attempting to UI*. To solve AP, the ‘practice of inferring’ that Harry must adopt is one of inferring ‘in accordance with the UI principle’ *in appropriate circumstances*. The stated instructions do not tell Harry how to adopt that practice.

Suppose that Harry decides to follow the instruction as stated. So, he tells himself, time and again, to infer an instance from any universal he believes. He realizes from the start that he won’t be following this instruction all the time: he has other things to do in life, like drinking beer and answering his emails. Still, he thinks that if he follows this instruction *often enough* he will establish a UI skill. Suppose he does establish a skill of sorts. Would that solve the AP? Well, maybe it would solve *some* version of AP, but not an interesting version. For, solving an interesting version demands showing how, on the basis of knowledge of the UI principle, the UI rule can become a part of Harry’s ‘evidential system’ for belief

formation, *just as it is a part of the system of normal rational adults*.¹⁷ Yet what Harry may have learnt is a skill of ‘mindlessly’ inferring a mass of information that is useless to him: he takes a random belief about *F*s in general and infers something about a random *F* in particular, whether or not he has any interest in that *F*; he infers by UI when there is no point in doing so and does not infer by UI when he really needs to. This ‘skill’ is very far from the UI skill of normal rational adults. Acquiring it does not provide what an interesting solution to AP should provide.

The moral of this? Harry needs to understand that *if* he is in the appropriate sort of circumstances *C*, *then* he follows the instruction. In other words, he would need to behave as if he had in mind a *conditional imperative*, of the form, “if *C*, do *A*!” Such a conditional imperative stands in contrast with a simple one of the form ‘Do *A*!’, exemplified by, ‘Shut the door!’, or, more pertinently, our stated instruction above. The problem for our neat solution is that following a conditional imperative seems to require an inference *that looks very like MP*.¹⁸ Worse, because this imperative must generalize the appropriate circumstances to UI, following it seems to require using an inference that *looks very like UI*. Harry has to put together his recognition of condition *C*—that he has a belief about all *F*s and wants to know about a particular *F*—and infer from the conditional imperative that he should therefore do *A*—infer an instance from a universal about *F*s. In Harry’s case, the reasoning process would look like this:

Argument C

- (C1) If I want to find out about a certain entity *x*, I believe that *x* is an *F*, and I believe all *F*s are *G*, then infer that *x* is *G*.
- (C2) I want to find out about a certain entity, Alex, I believe that Alex is an animal in the movie, and I believe all animals in the movie talk.
- (C3) So, infer Alex talks.

Is this reasoning process *really* an inference? Some have doubted that arguments which have imperative statements as premises or conclusions are inferences at all.¹⁹ If that were right, perhaps we could take AP as solved. But, we think Argument C has the key characteristics of an inference: (C1) and (C2) are reasons for doing (C3); more strongly, believing (C1) and (C2) but refusing to conclude (C3)

¹⁷ This reminds us of the Harman-inspired point in note 4. For the UI rule to be part of our epistemic system, we have to reason by it only when appropriate. Furthermore, having that rule involves more than moving from *beliefs* as premises to *beliefs* as conclusion when appropriate: it’s a matter of seeing what follows from premises when appropriate, whatever you believe. For the sake of simplicity, we continue to frame UI as an inference from beliefs to beliefs, though this results in a simplified picture of a person’s actual inferential life.

¹⁸ Following a conditional imperative (without instantiating) seems to require this sort of inference:

- (D1) If *C*, do *A*!
- (D2) *C*
- (D3) So, do *A*!

Our discussion in the following paragraphs is applicable to this MP-like argument, *mutatis mutandis*, to defend the position that it is an inference.

¹⁹ Bernard A.O. Williams is a notable example (Williams 1963, pp. 30-36).

seems like a logical mistake. Not only is (C) an inference, it seems to be a basic one closely related to UI and MP. How could Harry adopt *it*? So, rather than having solved AP, we have, at best, simply moved its focus to another inference.

In sum, Harry can't train himself to infer by UI from a general and conditional imperative because that would require that he was already able to infer by a basic inference close to MP and UI. And we saw earlier that he cannot train himself using a simple one. It is a bit surprising that Harry cannot solve AP by self-training. Still, we wonder just how interesting this is given that there is another way for him to rationally acquire the UI skill.

8. A Quinean coach-training answer

A good way for any person to gain a skill is to *hire a coach*. They both have declarative knowledge of a skill—they accept its 'principle'—which yields simple instructions. Obviously, the coach knows when it is appropriate to deliver the instructions.²⁰ Thus, consider Blanka who wants to learn the Fosbury Flop. In an appropriate setting she will be told by her athletics coach to 'approach the bar at your greatest controlled speed' and so on. After many repetitions she will have acquired the skill. Then, whenever she is in an appropriate circumstance and wants to high jump, she will execute a Fosbury Flop. And, we emphasize, this is a *rational* way for her to move from the principle to the skill. She has picked up the skill 'on the basis of', 'in virtue of', accepting the principle of the Fosbury Flop.

Analogously, suppose Harry comes to accept the UI principle. Then, it is rational of him to hire an inference coach to help him train to become competent in this difficult new skill of inferring by the UI rule. Harry's coach starts him off with the following inference.

1. All the animals in the movie talk.
2. Alex is an animal in the movie.
3. So, Alex talks.

The coach presents Harry with 1) and 2) in the appropriate circumstances; that is, Harry believes 1) and 2) and wants information about Alex. Then the coach gives Harry a simple imperative: 'infer that Alex talks.' Suppose that Harry obeys: he has come to believe that Alex the lion talks. It is not yet appropriate, of course, to say that Harry is governed by the UI rule; he's just done as he was told. But the inference coach then takes Harry through numerous such examples. In each case, the coach starts by presenting Harry with two premises of the forms 'All F s are G s' and ' $F(a)$ ' in circumstances in which Harry believes them and wants information about the object in question. The coach then instructs Harry to infer something of the form ' $G(a)$ '. After many repetitions we can expect Harry to become *disposed* to infer something of the form ' $G(a)$ ' in these circumstances.²¹ His reasoning is now governed by UI. Then,

²⁰ We are not concerned, of course, with how the coach knows all this: he might have got it from another coach, but he might have got it by implicit learning (§4) or partly innately.

²¹ Note that Harry is being taught to *infer* something of the form ' $G(a)$ ' from the premises not just *associate* ' $G(a)$ ' with them; he is being taught to treat the premises as a *reason* for ' $G(a)$ '. We assume Harry has been making other inferences before these lessons, albeit not by UI. He can, of course, learn to UI from the coach's instruction to 'infer that $G(a)$ ' *only if* he understands the instruction to mean that he

whenever Harry is in an appropriate circumstance wanting information about an object, he will execute a UI inference. And, we emphasize, the process by which Harry has put himself in this happy situation is as *rational* as the way that Blanka put herself in hers. Harry has picked up the skill ‘on the basis of’, ‘in virtue of’, accepting the UI principle. In both cases, acceptance of a principle leads to a rational course of action that brings about a skill. There is no cause to demand more ‘rationality’ in learning to reason than in learning any other skill. This Quinean coach-training answer illustrates how Harry could adopt UI, but such training could be applied, *mutatis mutandis*, to a person learning the other basic logical principles, such as MP, that concern AP.

Objection 1. ‘Even if the *decision* to follow the coach’s instructions is rational, the *process* of following them is not rational’. So, this proposed solution cheats by sneaking in an arational process.’ But, as we illustrated in §4, psychology (and commonsense) shows that skill learning *is* a largely arational process. To the extent that our skill learning is *explicit*—starting with a bit of declarative knowledge—the rational parts of the process are, first, in deciding to learn the skill described; and, second, in using the description to guide the execution of the decision. Explicit skill learning is an intentional activity but what we intentionally set in motion is an arational process; that is an empirical fact. We cannot exclude this process and still have an interesting answer to the Adoption Question.

Objection 2. But, as Padró says, ‘acceptance of the principle should *guide* Harry’s UI inferential practice’ (##4) and that is not what your training has achieved. The objection reflects an over-intellectualizing of inference that we resisted earlier (§3). We know already that skills need not involve any central processor representation of their ‘principles’. It remains an empirical issue whether skills even involve those representations in a sub-central module. If they do, then, *in that sense*, the content of the UI principle will indeed guide Harry’s UI-ing, because Harry has acquired the UI skill. If they do not, as we predict, then that content need no more guide Harry’s UI-ing than it guides anyone else’s.

Our Quinean coach-training answer is an interesting solution to AP because it reflects how skill acquisition, to the extent that it is rational, *actually* occurs. For that reason, it is also a scientifically respectable *naturalistic* answer. However, it is not the only solution to AP.

Consider another scenario for Harry. Suppose Harry accepts the UI principle and decides to undergo brain surgery to acquire the skill of inferring by UI. We earlier dismissed as ‘boring’ the idea of coming to use a rule ‘by mere surgery’ (§2). The present scenario is *not quite* that because Harry has made a *rational* decision for surgery based on his acceptance of the UI principle. Still, it seems a far less interesting answer to AP than the coach-training answer. Why? Because, adoption by surgery is not a matter of learning and is entirely non-cognitive (after the decision), whereas adoption by coach-training is a matter of learning and, though largely arational, is entirely cognitive: Harry is put through a cognitive skill-learning process (just as someone might be to learn chess).

This question about the interest of solutions to AP brings us back to the question of the interest of AP itself. AP is surely interesting as a challenge to provide a solution *using the devices of skill acquisition that organisms actually have*. It is surely uninteresting if those actual devices are excluded from any possible solution from the beginning.

should do the sort of thing that he does when he infers rather than merely associates or, for that matter, Fosbury Flops. If he lacked that understanding then he would be unteachable. But that is not interesting.

Indeed, an actual skill-learning process yields a solution to AP: it is *in principle* possible for someone who does not reason by UI to come to do so as a result of accepting the UI principle, if that acceptance results in hiring an inference coach to train. It is, of course, very hard to imagine someone *actually* doing so. Why? There are, after all, many organisms that do not reason by UI. Can't we imagine teaching them to do so in the way outlined? Well, the problem is that we have to imagine an organism that is sufficiently advanced to grasp the UI *principle*, something that humans managed only recently, and yet not advanced enough to be already governed by the UI *rule*, a fundamental rule of reasoning, something that our hominid ancestors probably achieved millions of years ago. That isn't easy to imagine! Indeed, perhaps Quine and Kripke can have another Kumbaya moment here. Perhaps if there were any organism, whether human or not, that had never inferred by UI, it would not be psychologically possible for that organism to adopt the UI rule as a result of accepting the UI principle. But we don't feel surprised or bothered by this. Should we be?

9. Conclusion

Can we change our logic and if so how? According to Kripke and Padró's Adoption Problem (AP), one sort of change is impossible: we cannot, on the basis of accepting a basic logical principle like UI, adopt the practice of reasoning by UI. We proposed a naturalistic Quinean response.

There are two questions about changing our logic. First, the **Acceptance Question**. It is obvious that we can accept a new basic logical principle, because logicians sometimes have over the last 3000 years. But how is it *rational* to do so? From our Quinean perspective, we cannot appeal to a priori knowledge to answer this notoriously difficult question: the answer must be empirical. We offered some suggestions as to how this might go, including appeals to Quine's metaphor of Neurath's boat. Our answer, like any answer, faced the problems of 'rule circularity' and 'self-defeat'. In considering the Acceptance Question, we acknowledged to Kripke that logical principles are indeed different from other propositions both in the way they affect others and the way others, and experience, affect them.

Second, the **Adoption Question**. Can a person, on the basis of accepting a new basic logical principle, rationally change her inferential practices? Our empirical hypothesis is that theoretical changes in our logic have caused lots of changes in our practices. The hypothesis does not seem bold for non-deductive ampliative reasoning. Furthermore, we argued, it is plausible that some changes in our practices have resulted from our acceptance of deductive principles. But what about AP? Could someone who did not already infer by a basic logical principle, like that of UI, come to do so on the basis of accepting its principle?

We have argued that, in principle, a person could. Our argument relied on empirically-supported claims from the psychology of skills to support our views of what it is for our reasoning to be governed by rules and how we can rationally come to be so governed. 'Declarative knowledge' of a skill, whether the skill of jumping by the Fosbury Flop or inferring by a basic logical principle, can help acquire the 'procedural knowledge' of the skill only by yielding the instructions used in 'explicit learning' to train someone in the skill. But though a person could train herself to Fosbury Flop, we have seen, surprisingly, that she could not *train herself* to infer in accordance with a basic logical principle; to that extent the AP is real. However, just as it is rational to hire an athletics coach to train to Fosbury Flop, it is rational to

hire an inference coach to train to infer. In this way, a person could rationally change her inferential practices on the basis of accepting a basic logical principle.

Right through, we have resisted over-intellectualizing inference. Reasoning, like Fosbury Flopping, is an intentional activity, but what we intentionally set in motion is simply a process governed by sub-centrally embodied rules (whether represented or not). And what we rationally set in motion when we train to Fosbury Flop or to infer by a new rule is largely not a rational process. We know from science (and commonsense) that such training is the only way to use knowledge of principles to master a skill. So all attempts to exclude such training from a solution to AP render AP uninteresting and so should be resisted.

ACKNOWLEDGEMENTS

We are indebted to Daniel Boyd, Saul Kripke, Romina Padró, Crispin Wright, and two anonymous referees for comments on a draft.

REFERENCES

- Anderson, John R. 1980, *Cognitive Psychology and its Implications* (San Francisco: WH Freeman)
- Boghossian, Paul 2000, 'Knowledge of Logic', in *New Essays on the A Priori*, eds. Paul Boghossian and Christopher Peacocke (Oxford: Oxford University Press)
- 2001, 'How are Objective Epistemic Reasons Possible?', in *Philosophical Studies* 106
- Burgin, Mark 2016, *Theory of Knowledge: Structures and Processes* (New Jersey: World Scientific)
- Carroll, Lewis 1895, 'What the Tortoise Said to Achilles', in *Mind* 104
- Devitt, Michael 2005, 'There Is No A Priori', in *Contemporary Debates in Epistemology*, eds. Matthias Steup, John Turri, and Ernest Sosa (Chichester: Wiley-Blackwell)
- 2006, *Ignorance of Language* (Oxford: Oxford University Press)
- 2011, 'No Place for the A Priori', in *What Place for the A Priori?*, eds. Michael J. Shaffer and Michael L. Veber (Chicago and La Salle: Open Court Publishing Company)
- Field, Hartry 1996, 'The A Prioricity of Logic', in *Proceedings of the Aristotelian Society* 96
- 1998, 'Epistemological Nonfactualism and the A Prioricity of Logic', in *Philosophical Studies* 92
- Papineau, David 1993, *Philosophical Naturalism* (Oxford: Blackwell Publishers)
- Psillos, Stathis 1999, *Scientific Realism: How Science Tracks Truth* (London: Routledge)
- Quine, W. V. O. 1951, 'Two Dogmas of Empiricism', in *The Philosophical Review* 60
- Schacter, Daniel L. 1999, 'The Seven Sins of Memory: Insights From Psychology and Cognitive Neuroscience', in *American Psychologist* 54
- Sun, Ron 2003, 'Connectionist Implementationalism and Hybrid Systems', in *The Encyclopedia of Cognitive Science*, ed. Lynn Nadel (London: MacMillan)
- Sun, Ron, Merrill, Edward, and Peterson, Todd 2001, 'From implicit skills to explicit knowledge: a bottom-up model of skill learning', in *Cognitive science* 25
- Tenison, Caitlin and Anderson, John R. 2015, 'Modeling the Distinct Phases of Skill Acquisition', in *Journal of Experimental Psychology: Learning, Memory, and Cognition* 42
- Williams, Bernard AO and Geach, P. T. 1963, 'Imperative inference', in *Analysis* 23