Reasoned Change in Logic

Elijah Chudnoff

Abstract: By a reasoned change in logic I mean a change in the logic with which you make inferences that is based on your evidence. An argument sourced in recently published material Kripke lectured on in the 1970s, and dubbed the Adoption Problem by Birman (then Padró) in her 2015 dissertation, challenges the possibility of reasoned changes in logic. I explain why evidentialists should be alarmed by this challenge, and then I go on to dispel it. The Adoption Problem rests on a failure to distinguish between logical principles such as Universal Instantiation and Modus Ponens which might or might not govern your inferences with superficially similar laws which must govern your cognitive architecture.

Keywords: inference, basing, evidentialism, adoption problem, cognitive architecture, logical evidence, revising logic

Is your logic responsive to your evidence?

The meaning of this question is not immediately obvious. By your logic I mean a subset of the inference rules you follow when you make inferences. In recently revived scholastic jargon, I'm talking about your *logica utens*, not your *logica docens*.[1] Example rules include Universal Instantiation (UI) and Modus Ponens (MP):

> (UI) For any A and t, if 'For all x, A' is a premise, then it is correct to infer 'A(t/x)'.

> (MP) For any A and B, if 'A' and 'If A, then B' are premises, then it is correct to infer 'B'.

These exhibit a basicness and generality that distinguish them from inference rules encoding empirical regularities, such as:

> (Rain) For many t, if "The streets are wet at t' is a premise, then it is correct to infer 'It rained shortly before t'.

By your evidence I mean considerations of yours that count for or against your having attitudes such as belief and disbelief.[2] The observations you've made, the testimony you've received, the intuitions you've had, and the reasoning you've engaged in all comprise your evidence. Responsiveness is a trickier notion. In this paper, I work with the following simple

---

[1] Its recent popularity in the literature on revising logic seems to trace back to (Priest 2016). The distinction raises substantial interpretive issues elided here, but the basic idea is that your *logica docens* is the logic you preach while your *logica utens* is the logic you practice.

[2] I have in mind considerations that bear on the correctness of the relevant attitudes. Whether some such qualification should be added to the formulation in the text in order to rule out practical reasons for belief is an issue I set aside here.

conception: if your logic includes a rule R because your evidence leads you to believe that R is true or otherwise acceptable, then your logic is responsive to your evidence.[3]

Threadbare as the foregoing exposition can seem in light of the weightiness of the notions dealt with, it suffices to locate my question in relation to two familiar controversies. The first provides a useful analogy. According to Pylysyn (1999) and many others, your visual system, or a significant part of it, is cognitively impenetrable. To illustrate the idea, consider the case of known illusions, such as the Müller-Lyer illusion:

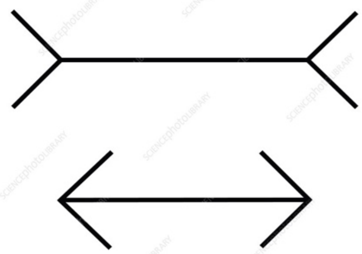

Measurement will lead you to believe the two lines are the same in length, but the top line will continue to look longer than the bottom line. Your visual system generates this look by following rules that are unresponsive to your evidence. My question is akin to asking whether your logic is similarly unresponsive.

The second controversy is one that substantively bears on how my question should be answered. I have in mind the problematic sourced in lectures Kripke gave in the 70s and, due to Birman's (then Padró's) 2015 dissertation, now known as the Adoption Problem: "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" (Birman 2023, 3). In Kripke's own words: "logic, even if one tries to throw intuitions to the wind, cannot be just like geometry because one cannot adopt the logical laws as hypotheses and draw the consequences. You need logic in order to draw these consequences" (Kripke 2023, 19). These claims suggest a negative, or perhaps a qualifiedly negative, answer to my question: your logic cannot be responsive, or as responsive, or similarly responsive, to your evidence as your beliefs are. If it could be, then adopting a logic shouldn't be more problematic than adopting a geometry or any other theory.

Aside from a natural interest in evidence's capacity to influence and rationalize our epistemic states and activities, evidentialists should care about the Adoption Problem. Evidentialism includes a theory of well-founded belief and logical inference is a way of forming well-founded beliefs. In Section 1, I'll suggest that plausible evidentialist ideas about well-founded belief demand that your logic be responsive to your evidence. The balance of the paper addresses the Adoption Problem. Section 2 presents the problem more fully than these introductory remarks; Section 3 elaborates on some ideas about our cognitive architecture; and Section 4 draws on those ideas to show how it is possible to adopt logical principles such as UI and MP. Section 5 concludes.

---

[3] I've given a sufficient condition for responsiveness. It might be that there are weaker sufficient conditions, but the one in the text serves present purposes since it corresponds to how contributors to the literature on the Adoption Problem understand what it is to adopt a logic.

1. Evidence and Logical Inference

As I understand it, inference in general is agent-level epistemic basing. When you make an inference, you assume responsibility for basing an attitude such as belief on some considerations such as other beliefs. In this paper I focus on a special class of inferences I call deliberate inferences. These have the following structure:

Inferential rule[4] $-_{\text{generates}}\rightarrow$ inferential seeming $-_{\text{guides}}\rightarrow$ inferential act

By an inferential seeming I mean an experience as of some conclusion following from some premises. Compare the following arguments:

(A) If time is money, then cash is king.
(B) Time is money.
(C) So, cash is king.

(D) If if time is money, then cash is king, then time is money.
(E) So, time is money.

If you are like me, then when you go through the first argument premises (A) and (B) are experienced as supporting conclusion (C), but when you go through the second argument premise (D) is not experienced as supporting conclusion (E), though you might recognize that (E) does follow from (D) by Peirce's Law. In the first case there is an inferential seeming, and in the second case there is not. Inferential seemings are generated by the general inference rules you follow, and they in turn guide the particular inferences you make. If the process only involves a logical inference rule, then you've made a logical inference.

Evidentialism includes a theory of well-founded belief:

WF S's doxastic attitude D at t toward proposition p is well-founded if and only if
(i) having D toward p is justified for S at t; and
(ii) S has D toward p on the basis of some body of evidence e, such that
 (a) S has e as evidence at t;
 (b) having D toward p fits e; and
 (c) there is no more inclusive body of evidence e' had by S at t such that having D toward p does not fit e'.[5]

Logical inference is a way of establishing the sort of basing relation cited in WF (ii), so evidentialists should be interested in how it works. And there are reasons to think there are conditions in play that WF does not include or at least does not make explicit.

Consider the following argument:

---

[4] Since rules themselves are abstract, what's meant here is an appropriate mental representation of the inferential rule.
[5] (Conee and Feldman 2004, 93; a reprint of Feldman and Conee 1985).

(F) If time is money, then it isn't the case that time isn't money.
(G) It isn't the case that time isn't money.
(H) So, time is money.

Imagine someone whose evidence includes (F) and (G) and who experiences (F) and (G) as supporting (H) because they follow the bad rule Modus Morons (MM):

> (MM) For any A and B, if 'B' and 'If A, then B' are premises, then it is correct to infer 'A'.[6]

(F) and (G) justify believing (H) because (G) alone justifies believing (H), but arguably this person's belief is not well-founded. There is something epistemically amiss about it.[7] Here is a natural story about what that is.

Inferring (H) from (F) and (G) by (MM) fails to produce a well-founded belief in (H) because, unlike (MP), (MM) is an inappropriate inference rule. What makes (MP) appropriate and (MM) inappropriate?

We can observe that (MP) licenses logically valid inferences and (MM) licenses logically invalid inferences. This is one difference between them. However, for the purposes of tracking well-founded belief, there is reason to think that some logically valid inferences shouldn't be licensed for some thinkers because the validity of those inferences is opaque to those thinkers. For example, a rule licensing inferences from the Peano Axioms to complex number theoretic truths would not be appropriate for most thinkers. The relevant sort of appropriateness and inappropriateness should account for the difference between this rule and (MP), which is appropriate for most thinkers.

A more plausible idea is that appropriate rules are appropriate because they are justified, and inappropriate rules are inappropriate because they are unjustified. Suppose something along these lines is correct and evidentialism is the correct theory of justification. Then well-founded beliefs should meet the following condition: if the basing relation cited in WF (ii) is established by inference, then the rule followed in making that inference should be one supported by your evidence. It is a short step to the idea that your logic should be responsive to your evidence since the rule shouldn't just be supported by your evidence but should be followed because of that support. Hence evidentialists should care about the Adoption Problem, to which I now turn.

2. The Adoption Problem

---

[6] The fallacy is affirming the consequent; "Modus Morons" comes from (Haack 1976).

[7] The case is similar to McCain's (2014, 95) "Bad Reasoning" example. There he suggests the case is one in which the basing relation fails to be established at all. (McCain 2023) develops a view somewhat closer to mine, though he focuses on cases in which basing fails to establish well-foundedness because of what I'd call a problematic inferential seeming rather than a problematic inference rule.

To adopt a logical principle is to begin following it because you have accepted it (cf. Birman 2023, 3 – 4). New inferential practice must be guided by the content of new logical belief.[8] Consider a case that hasn't generated controversy, such as Modus Tollens (MT):

> (MT) For any A and B, if 'not-B' and 'If A, then B' are premises, then it is correct to infer 'not-A'.

Suppose a logical novice initially doesn't follow (MT), is then shown how for any A and B, 'not-B' and 'If A, then B' logically imply 'not-A,' and thereby comes to accept that (MT) is true. Now they are given the premises, 'Cash isn't king,' and 'If time is money, then cash is king.' With some effort they recognize that these premises and (MT) make it correct to infer, 'Time isn't money.' So, they go ahead with this inference. They have adopted (MT).

Modus Tollens has not generated controversy because it is not basic. Arguably, if we zoom into the scenario just described, then we will find that the novice is able to follow (MT) because they already follow other logical principles such as (UI) and (MP). It is these basic logical principles that are supposed to be unadoptable. Here is how the reasoning goes for (UI) (cf. Kripke 2023, 15; Birman 2023, 3 - 7). Suppose someone does not follow (UI), for example they do not infer (2) from (1):

(1) All ravens are black.
(2) This raven is black.

In order to begin following (UI) in this case because of having accepted (UI), they must infer (5) from (3) and (4):

(3) For any A and t, if 'For all x, A' is a premise, then it is correct to infer 'A(t/x)'.
(4) (1) is a premise, and if (1) is a premise, then it is correct to infer (2).
(5) It is correct to infer (2).

It is worth calibrating the process described here with the picture of deliberate inference sketched in the previous section.

Inferential rule $\xrightarrow{\text{generates}}$ inferential seeming $\xrightarrow{\text{guides}}$ inferential act
(3) (4) (5) (1)/(2)

(3) is the content of an inferential rule, (4) describes how the rule generates an inferential seeming, (5) is the content of that inferential seeming, and (1)/(2) is the inference made in response to it. The question of how guidance works is left open, and I will continue to bracket it in this paper.[9]

Now the problem is that if generation requires an inference from (3) to (4), more specifically from (3) to the second conjunct of (4), then our imagined thinker already

---

[8] This is definitional of adoption.
[9] I address it in detail in (Chudnoff forthcoming).

follows (UI). So, (UI) cannot be adopted. If a subject follows it, e.g., when inferring (4) from (3), then no adoption is needed. And if the subject does not follow it, e.g., in failing to infer (4) from (3), then no adoption is possible. This is the Adoption Problem for (UI). Granting that the inference from (4) to (5) is also required, it is clear that similar reasoning applies to (MP).

The most contentious step is the one linking the (1)/(2) inference to the (3)&(4)/(5) inference. Some supporters of the Adoption Problem aim to justify it on the basis of more general truths about rule-following (Finn 2019). An alternative reaction is to argue that variants on the original Adoption Problem demonstrate similar a priori limits to adopting a logic without relying on the same contentious step (Boghossian and Wright 2023). Skeptics about the Adoption Problem reject the contentious step, and with it the idea that there are any a priori limits to adopting a logic (Devitt and Roberts 2023).

In my view, there is a connection between the (1)/(2) inference and a process with a (3)&(4)/5 structure, but that process is not an inferential process, and so there is no Adoption Problem. I reject what might be called the Inference Assumption:

> Inference Assumption: Generating a representation, e.g., an inferential seeming, of what general inference rules mandate in a particular case requires making an inference.

To give an alternative non-inferential account of the process we need to consider our cognitive architecture.

3. Cognitive Architecture

One enduring insight to have emerged from the cognitive sciences is that intelligent behavior is a function of shared and stable cognitive mechanisms operating on individualized and changing bodies of information and goals. The shared and stable cognitive mechanisms responsible for intelligent behavior in humans comprise our cognitive architecture.[10] Though the ideas I rely on are common ground among theories of

[10] This leaves room for refinement, but the basic idea is common across research programs associated with different proposals about the specifics. Here, for example, is John Laird in his book length presentation of the Soar cognitive architecture: “a cognitive architecture provides the fixed processes and memories and their associated algorithms and data structures to acquire, represent, and process knowledge about the environment and tasks for moment-to-moment reasoning, problem solving, and goal-oriented behavior” (Laird 2012, pg. 8). And in his most recent book length presentation of the ACT-R cognitive architecture, J. R. Anderson writes, “A cognitive architecture is a specification of the structure of the brain at a level of abstraction that explains how it achieves the function of the mind” (Anderson 2007, pg. 7). This takes a bit more pondering, but it is clear from the surrounding discussion that Anderson aims to capture the same basic idea as Laird. He approvingly quotes Allen Newell’s formulation in *Unified Theories of Cognition* according to which human cognitive architecture consists of “the fixed (or slowly varying) structure that forms the framework for the immediate processes of cognitive performance and learning” (Newell 1990, pg. 111).

cognitive architecture, it will be convenient to explain them in the context of a specific model. I'll consider a version of ACT-R, diagramed below.[11]

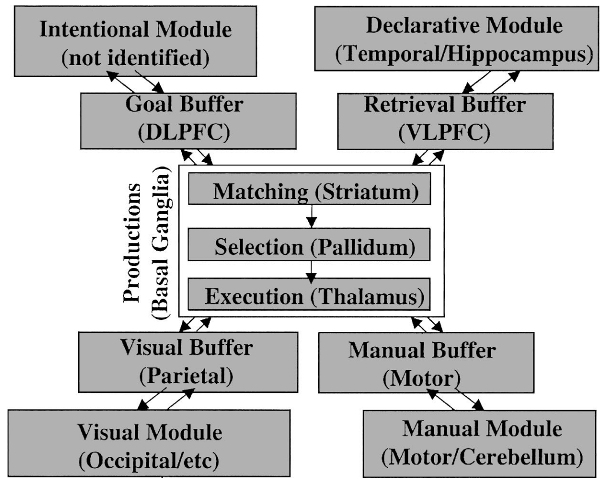

ACT-R

The diagram represents hypotheses about neural implementation, but for present purposes what matters are the three kinds of mental structure and their interaction.

There are modules, their buffers, and a production system. Modules store and process proprietary content, such as visual information (visual), motor programs (manual), goals and subgoals (intentional), and associatively linked facts (declarative). Limited amounts of this content are transferred into and out of the modules in the service of performing cognitive tasks; the buffers are temporary stores for this content in motion. The production system is the engine. It cycles through a process of matching the overall state of the buffers at one moment to conditions for making changes to the overall state of the buffers at the next moment. If conditions for making incompatible changes are met, then the one currently associated with the highest utility is selected for execution.[12]

This is a bare bones description and doesn't begin to suggest the explanatory power of and experimental support for understanding the mind as organized along such lines. Interested readers should consult the referenced literature. What matters for present purposes is how our cognitive architecture processes rules.

[11] The version of ACT-R I consider is set out in (Anderson et al 2004). Later versions add to the model without rejecting the components I discuss, and the earlier version is adequate for present purposes. (Anderson 2007) presents ACT-R in a way that highlights its embodiment of agreed on ideas about cognitive architecture. (Laird et al 2017) compares some of its details with those of other leading models while distilling broad points of convergence. If you look outside the literature dedicated to developing computer models, then the terminology changes, but the overall shape attributed to cognitive architectures is similar. See for example the cognitive architecture attributed to honeybees in (Menzel and Giurfa 2001).

[12] The utility of a production is determined by factors such as the value of the current goal, the estimated probability that the production will succeed, and the estimated cost of selecting it. ACT-R includes equations for these and related quantities, but they will not figure in the present discussion.

The production system maps patterns (contents of the buffers) to actions (changes to the buffers). The mappings can be represented as rules, called production rules or productions. Here is an example adapted from a list of productions for multidigit subtraction (Anderson 2007, 153 – 154):

Condition	Action
If the goal is to process a column, *And* the top digit is not smaller than the bottom digit	*Then* set the goal to finding the difference between digits, *And* retrieve their difference

The condition tests the contents of the goal buffer and the visual buffer. If it is met, then two actions are performed. The goal is set to finding the difference between digits, and a request for that information with respect to the digits in the column being processed is sent to the retrieval buffer.[13] If a person were ascertaining what a rule along these lines mandates when starting on the problem, 24 – 12, then they might use (UI) and (MP) to reason as follows:

(1) If my goal is to process a column, and the top digit is not smaller than the bottom digit, then I should set my goal to finding the difference between digits, and I should retrieve that fact for these digits.
(2) My goal is to process 4 – 2, and 4 is not smaller than 2, and if my goal is to process 4 – 2, and 4 is not smaller than 2, then I should set my goal to finding the difference between digits, and I should retrieve that fact for 4 and 2.
(3) I should set my goal to finding the difference between digits, and I should retrieve that fact for 4 and 2.

We should not conclude that a person must follow anything like (UI) or (MP) in order for the production under consideration to partly explain their ability to do subtraction problems.

First, production systems are not persons, nor do they make up the whole of a cognitive architecture. So, whatever the production system that might be part of your cognitive architecture is doing, it need not be something that you are doing.

Second, production systems themselves likely do not follow or even conform to (UI) or (MP). When we represent productions as condition-action pairs in English they appear to be generalized conditionals, and so apt to occur in inferences governed by (UI) and (MP). But the appearance is merely an artifact of this particular, and at best heuristic, way of representing productions. To know what a production system is doing we need to know how the production system itself represents and processes the productions on which its match-select-execute cycle operates. Given that the cycle includes a selection phase, we know that the action-parts of productions are not detached from the condition-parts in a way that mirrors (MP). Further, the condition-parts plausibly consist of templates for patterns in sets of buffers. Templates match or fail to match; but unlike antecedents of conditionals,

[13] In this example, I'm supposing that the visual module computes the relative size of the digits, and I'm ignoring the selection phase in the production cycle.

they are not true or false. The action-parts plausibly consist of sets of operators. Operators are executed or not executed; but unlike consequents of conditionals, they are not true or false.

A final point, congruent with these observations, is that production systems can be implemented in devices and creatures that are incapable of propositional thought.[14] On their own, production systems are simple mechanisms; it is only in the context of a suitable cognitive architecture that they constitute engines capable of generating intelligent behavior.

So, there is no reason to think that if intelligent behavior depends on production rules, then it must also depend on inference rules such as (UI) or (MP). From the supposition that your intelligent behavior derives from a cognitive architecture along the lines described by ACT-R and its companions in the literature, most of which incorporate production systems, it does not follow that you or any parts of you follow or even conform to (UI) or (MP). It is a further question whether the same supposition provides material for explaining how you might begin to follow (UI) and (MP) by accepting them, i.e., how you might adopt (UI) and (MP). I turn to this question in the next section.

4. The Architecture of Logic Adoption

When they first confront multidigit subtraction problems, children are not already equipped with the production discussed in the previous section. The correct production must be learned, and typically this occurs through instruction and example. The focus here will be instruction. A simplified ACT-R account of how children might learn our example production will suggest a model for how adults, or children, might adopt logical principles.

To learn from instructions, children must be able to comprehend and follow them. The instruction in this case can be expressed as follows:

> When processing a column, if the top digit is not smaller than the bottom digit, then the correct operation is to find their difference.

Comprehending the instruction consists in storing an appropriate representation of it in declarative memory. The representation is indexed to certain retrieval requests, and it returns certain operators when those requests are made. Following the instruction requires productions. Some productions make retrieval requests to which the instruction is indexed, and other productions execute the instruction's operators. The following will suffice for illustrative purposes:

Condition	Action
If the goal is to process a column, *And* the top digit is not smaller than the bottom digit	*Then* retrieve an instruction for that kind of column

[14] Though Menzel and Giurfa (2001) do not use the term "production system," the cognitive architecture that they attribute to honeybees performs the same functions.

Condition	Action
If the retrieved instruction is to find the difference between two digits	*Then* set the goal to finding the difference between digits, *And* retrieve their difference

Consider a child who receives the instruction, has the abilities required to follow it, and is starting on the problem, 24 – 12. The first production leads them to recall the instruction for columns in which the top digit is not smaller than the bottom digit. Once the instruction is recalled, the second production leads them to recall the difference between 4 and 2. Their behavior will be similar to, though less efficient than, the behavior of the practiced subtractor who has the single production from the previous section.

The ACT-R learning mechanism called production compilation takes two productions of the sort described in this section as input and gives a single production of the sort described in the previous section as output. Specifically, the action of the first production results in the condition for the second production being satisfied, and production compilation cuts out the mediating call to the declarative module. It operates automatically whenever two productions standing in this sort of relation are executed. However, the single output production only begins to be selected for execution over the two input productions when its associated utility overtakes theirs due to its greater efficiency.[15] This is one reason performance speeds up with practice. So, altogether, we've seen that typical learning from instruction combines at least three more basic forms of learning: declarative learning of an instruction, procedural learning of a production, and gradual improvements to performance through updates in utilities.

A similar story can be told about adopting logical principles. Storing a suitable representation of (UI) in declarative memory suffices for accepting it. Let's suppose instruction or other evidence results in a representation of (UI) stored in declarative memory in something like the form reproduced here:

(UI) For any A and t, if 'For all x, A' is a premise, then it is correct to infer 'A(t/x)'.

Adopting (UI) requires following it because of having accepted it, and that requires productions for following inference rules. The subtraction example suggests the following:

Condition	Action
If the goal is to evaluate an inference, *And* the premise is 'For all x, A' and the conclusion is 'A(t/x)' for some A and t	*Then* retrieve an inference rule for that kind of inference

[15] This brief account abstracts from details about how utilities are updated. The basic idea combines noise and reinforcement learning. Noise in the system sometimes results in productions with lower utilities being selected over productions with higher utilities. Reinforcement learning adjusts the utility associated with the selected production in light of its effects. If it is more efficient than its rivals, then over time it will become the production with the higher utility.

Condition	Action
If a retrieved inference rule classifies an evaluated inference as correct	*Then* generate an inferential seeming of as of its conclusion following from its premises

Consider someone who has just accepted (UI), is equipped with these productions for following inference rules, and is now evaluating the inference from (1) to (2):

(1) All ravens are black.
(2) This raven is black.

The first production leads them to retrieve the inference rule for inferences of the form 'For all x, A'/'A(t/x)' for some A and t. It does so because (1)/(2) is of this form, so the production's condition is satisfied. There is no inference here, just matching. Once the rule is retrieved, the second production leads them to experience the inference from (1) to (2) as correct. Again, there is no inference here, just execution of an operator. So, their mind generates an inferential seeming of (2) as following from (1) without them or any part of them following or conforming to (UI) or (MP).[16] Suppose you go ahead with the inference that appears correct, i.e., you infer (2) from (1). Then you count as inferring (2) from (1) by following (UI). You follow (UI) in making the inference from (1) to (2), but you do not follow (UI) in generating the inferential seeming as of that inference being correct.

Backtracking a bit, as soon as the two productions for following (UI) fire in sequence, production compilation will result in the following new production:

Condition	Action
If the goal is to evaluate an inference, *And* the premise is 'For all x, A' and the conclusion is 'A(t/x)' for some A and t	*Then* generate an inferential seeming of as of its conclusion following from its premises

Initially it will not be selected for execution because it will start with a default low utility, but over time its greater efficiency will be calculated into the utility assigned to it, and it will be selected in preference to the two productions that require a call to the declarative module. I'm inclined to say that (UI) has already been adopted prior to this improvement in performance, but nothing in the present account rules out stricter proficiency standards on what counts as having adopted a logical principle.

So, there is a route from newly accepted general inference rules to particular inferential seemings that does not require making an inference. The Inference Assumption—that generating a representation, e.g., an inferential seeming, of what general

[16] Here I disagree with Finn's (2019) claim that (MP) and (UI) are unadoptable because they are self-governing. In the previous section I gave reasons for distinguishing the match-select-execute cycle from making inferences in accordance with (MP) and (UI). Once this distinction is made, then there is no reason to think (MP) and (UI) "govern such basic and fundamental patterns of inference that they underwrite the application of any logical rule, including themselves." (Finn 2019, 239).

inference rules mandate in a particular case requires making an inference—is false. And if it is false, then there is no Adoption Problem.[17]

5. Conclusion

An upshot of the foregoing is that while logic might be basic to inference, inference is not basic to cognition—including cognition that is responsive to evidence.[18] This opens the door to reasoned changes in logic. Whether there are sufficient reasons to walk through that door is a further question not addressed here.

Bibliography

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological review*, *111*(4), 1036.

Anderson, J. R. (2007). *How can the human mind occur in the physical universe?*. Oxford University Press.

Birman, R. (2023). The Adoption Problem and the Epistemology of Logic. *Mind*, online access.

Boghossian, P. and Wright, C. (2023). Kripke, Quine, the "Adoption Problem" and the Empirical Conception of Logic. *Mind*, online access.

Chudnoff, E. (forthcoming). Inferential Seemings. In *Oxford Studies in Philosophy of Mind 4*. (ed. Kriegel, U.). Oxford University Press.

Conee, E., & Feldman, R. (2004). *Evidentialism: Essays in epistemology*. Clarendon Press.

Devitt, M., & Roberts, J. R. (2023). Changing Our Logic: A Quinean Perspective. *Mind*, online access.

Feldman, R., and Conee, E. Evidentialism. *Philosophical Studies* 48 (1985): 15-34.

Finn, Suki (2019). The Adoption Problem and Anti-Exceptionalism about Logic. Australasian Journal of Logic 16 (7):231.

Haack, S. (1976), The Justification of Deduction. *Mind* 85, 112-119.

McCain, K. (2014). *Evidentialism and epistemic justification*. Routledge.

McCain, K. (2023). Evidentialism, Judgment, and Suspension: Meeting Sosa's Challenges. *Episteme*, 1-12.

Menzel, R., & Giurfa, M. (2001). Cognitive architecture of a mini-brain: the honeybee. *Trends in cognitive sciences*, *5*(2), 62-71.

Laird, J. E. (2019). *The Soar cognitive architecture*. MIT press.

---

[17] (Boghossian and Wright 2023, 11) criticize what they call the "Original Adoption Problem" on similar grounds. In light of the criticism, however, they suggest that general rules can influence particular acts via "Immediate Contentual Regulation." I am not persuaded that this is possible, and I believe there is need for intermediate states such as inferential seemings.

[18] (Devitt and Roberts 2023) also draw from the literature on cognitive architecture in order to address the Adoption Problem, but they focus on the role of training rather than responsiveness to evidence, hence Birman's (2023, 7) quick dismissal of their proposal. A similar dismissal would not apply to the account developed here.

Laird, J. E., Lebiere, C., & Rosenbloom, P. S. (2017). A standard model of the mind: Toward a common computational framework across artificial intelligence, cognitive science, neuroscience, and robotics. *Ai Magazine*, *38*(4), 13-26.

Newell, A. (1994). *Unified theories of cognition*. Harvard University Press.

Padró, R. (2015). *What the Tortoise Said to Kripke: the Adoption Problem and the Epistemology of Logic*. PhD thesis, CUNY.

Priest, G. (2014). Revising logic. In Rush, P., editor, The Metaphysics of Logic, chapter 12, pages 211–223. CUP, Cambridge.