Adopting an Inference Rule: A How-to Guide

William Nava*

February 22, 2025

Abstract

This paper introduces a novel account of the process by which agents adopt unfamiliar inference rules. On this account, adoption is a process during which agents are *inferentially guided* by an explicit statement of the rule they are adopting, but during which they do not use that rule. Rather, the ability to use the rule is the outcome of the process. This account avoids a regress objection to inferentially guided adoption recently posed by Boghossian and Wright. Adoption, on this model, requires the use of six privileged inference rules, including universal instantiation, modus ponens, and the naive truth rules. However, though these rules play a special role in adoption, they are not indispensable to the adoption process. On the contrary, it is possible to adopt without them *so long as one has restrictions of them*; and there are no minimum restrictions required for the adoption process to succeed. One upshot of this is that no familiar rule of logic is unadoptable.

1 Solving the adoption problem

In response to a number of longstanding problems, many philosophers and logicians have proposed subclassical logics. Some of these proposals are descriptive—i.e., claims to the effect that we already do reason subclassically. Many, however, are normative. They propose that we *adopt* a new logic in light of paradox or other theoretically recalcitrant phenomena.

Famously, Saul Kripke argued in the 1970's against the possibility of adopting a new logic. Kripke's argument,[1] which has come to the be known as the *adoption problem*, goes roughly as follows. Suppose Harry, who cannot antecedently instantiate, is told and accepts:

(Universals) All universal claims entail their instances.

Harry is also told and accepts a particular universal:

(Animals) All animals in the movie *Madagascar* talk.

*This paper is forthcoming in *Mind*. I extend a warm thanks to Hartry Field, Noga Gratvol, Marko Malink, Matthew Mandelkern, Graham Priest, Ali Rezaei, Crispin Wright, and an anonymous referee for useful comments and suggestions.

[1] See Kripke (2023). The case is from the canonical formulation by Birman (2015) (née Padró).

In order to infer an instance of the latter—say, 'Alex the lion talks'—Harry would need to first infer that (Animals) is an instance of (Universals). But this is by stipulation what he cannot do. Unless Harry can *already* instantiate, a statement expressing what instantiation is will not help. Therefore, Harry cannot adopt universal instantiation (UI).

To what rules might this kind of argument apply? It is by no means obvious that it is all or only logical ones. Intuitively, this argument type applies to rules that are, in general, used to inferentially move from a sentence that expresses a rule, along with some premises, to the corresponding conclusion. Following Birman (2015), call these rules *basic*.[2] Prima facie, if Harry-type arguments work, they show that exactly the basic rules are unadoptable.

Now, an immediate issue is that it's unclear how this argument affects proposals for *dropping*, as opposed to adding, rules. This is important because, as noted by Cohnitz and Nicolai (forthcoming), the vast majority of proposals for alternative logic are for dropping (or, more precisely, restricting) logical rules. If the argument were to apply only to proposals for adding rules, it would bear only on merely possible cases like that of Harry and not on live debates about logic.

However, the argument *does* bear on proposals for dropping rules—precisely the basic ones. The reason, in brief, is that rule adoption isn't some esoteric activity relevant only to deviant logicians. It is a commonplace practice of elementary importance to our cognitive lives. But, as we'll see, the Kripke-Birman argument turns on the assumption that basic rules are indispensable to rule adoption in general. For this reason, if their argument works, it also shows that one mustn't *drop* any basic rules, on pain of severe cognitive impairment.

There is a potentially more serious objection to the Kripke-Birman argument, raised by Boghossian and Wright (2023): its apparent appeal to the *inferential model of rule following*. This is the view that use of an inference rule requires one to make an inference from an explicit representation of that rule. The Harry argument does at least seem to turn on the assumption that if Harry is to use UI, he must do so by inferring from a sentence that expresses it. But this assumption generates a Lewis Carroll-type vicious regress, as convincingly argued by Boghossian (2014).

One may be tempted to say: so much the worse for the Kripke-Birman argument! But not so fast. It certainly seems that when we adopt new inference rules, we do so precisely under inferential guidance from statements of those rules. Indeed, it's not clear how else we might go about adopting new rules. So, if this objection successfully undermines the Kripke-Birman argument against adoption, it undermines the possibility of adoption, too!

My primary aim in this paper is to show that both adopting and dropping rules—even basic ones—is both possible and viable. This is so even if rule adoption requires inferential guidance from a statement of the rule, as it prima facie seems to. The key is to recognize that adoption is a temporally extended *process* of learning a new practice. Careful consideration of what this process involves reveals that one can learn to use a new rule, under inferential

[2]Birman introduces basic rules as those which are underived. However, as she rightly notes, which these are depends on one's derivation system. She therefore leaves the exact notion unspecified, stipulating that modus ponens and UI are basic. Her discussion, however, suggests that basic rules are just those to which Harry-type arguments seem to apply, i.e. those that are basic in my sense. See Birman (2015, fn. 56).

guidance from a statement of it, without this implying the inferential model of rule following; and without already having recourse to the basic rules.

By vindicating both the possibility of adopting basic rules, and the at least potential viability of dropping them, I defend the normative significance of major proposals for alternative logic. *Pace* Kripke and Birman, these proposals ask us to do something that is both possible and abductively evaluable. And, *pace* Boghossian and Wright, we can do it in just the way we seem to: under inferential guidance from statements of the rules in question.

It's worth flagging another feature of the account. As Kripke himself briefly hints at, and I will argue later, both naive introduction and elimination rules for a truth predicate are basic. But, given very minimal assumptions, these rules are inconsistent with classical logic. If the Krike-Birman picture of adoption is right, we should not drop any basic rules. It follows that we must drop non-basic classical rules, in order to avoid proofs of triviality using the truth rules—thereby adopting a subclassical logic. My account of adoption blocks this argument. Though the truth rules are indeed basic, adoption considerations do not bar us from dropping them for abductive reasons, as we might any other rule. So, my account at least makes room for the viability of classical logic.

The paper proceeds as follows. In the next section, I put forward the account of adoption as a learning process and address the regress objection. §3 models in more detail how the account works and draws out the list of basic rules. In §4, I show how the Kripke-Birman argument should be understood in light of my account—and why it ultimately fails. I then (§5) go on to draw some consequences for 'anti-exceptionalism' about logic. §6 discusses the conditions under which attempts at adoption might still fail on this account. In the final section (§7), I defuse a new adoption problem put forth by Boghossian and Wright (2023) and address a concern raised by them about the possibility of restricting a rule.

2 What is adoption?

'Adoption' is a term of art. Nevertheless, it picks out a perfectly quotidian cognitive process. My aim in this section is to draw out this familiar phenomenon and offer a definition of it that is compatible with how the term has been used in the existing literature.

It will help to begin by clarifying the target notion of inference. As I'll use the term, inference is an action that takes us from one mental state to another *in one cognitive step*. That is, if a movement from acceptance of some premises to acceptance of a conclusion is an inference, it does not involve any intervening mental states. It is a *direct* cognitive transition.[3]

Knowledge of a fact does not suffice for the ability to take such an action, even when the fact in some sense corresponds to an inference rule's content. Consider Irene, who is starting to learn how to play chess. Irene has been told, and remembers, the general rule for how a knight moves. Nevertheless, when given a position, she needs to think through how the rule applies in order to determine where she's allowed to move her knight. While Irene knows *that* knights move thus and so, she is not yet competent with the corresponding rule, because

[3]Of course, not *any* cognitive move of one step is an inference. This is only a necessary condition.

she cannot reliably move in one cognitive step from recognizing a chess position to what the possible knight moves are. She needs to make *multiple* inferences to reach the desired conclusion. This is in contrast to seasoned chess players, who do make that transition in one step and so can be said to reach the conclusion in *one* inference.

To be sure, we do sometimes call cognitive moves of multiple steps 'inferences'. But doing so in general would imply that the ability to follow a complicated and long train of reasoning—as in, say, making one's way through a non-trivial proof—constitutes competence with *an* inference rule. So as to maintain a tight association between inference rules and the actions they license, it is better to restrict the target notion of inference as indicated.

On to adoption itself. Adoption is a way of acquiring competence with the use of an unfamiliar inference rule. What distinguishes adoption from other kinds of rule acquisition is a certain kind of sensitivity to the rule's content.[4] As a first, blunt pass: when we adopt an inference rule, we acquire it *under guidance by an explicit representation* of that rule.[5]

In the vast majority of cases, our most general inference rules, including rules of logic, are either innate or learned by immersion in a social practice. Neither method involves guidance by an explicit representation of the rule's content. So our acquisition of these most general rules is not, at least typically, by adoption.

Even so, adoption as so far characterized is not at all uncommon. We often adopt inference rules as part of the process of deliberately acquiring a new concept. For example: learning to infer, from given chessboard positions, what the legal knight moves are, is part of the process of acquiring the concept KNIGHT. This inference rule is usually acquired via a process in which we are explicitly told what the rule is and are guided by our grasp of it.

Concept acquisition does not exhaust instances of adoption, however. My adopted ability to infer, from the lowering of a camera's f-stop,[6] that the image will have a shallower depth of field, postceded my grasp of the concepts F-STOP and DEPTH OF FIELD. Likewise, I adopted the chess rule *en passant* (which governs a somewhat uncommon situation) well after grasping the concept PAWN.

We ought not confuse the content-sensitivity distinctive of adoption with *deliberate* or *conscious* rule acquisition.[7] Imagine a case in which Pepper, who has no interest in developing expertise in photography, is faced with a series of photography tasks in which depth of field is important. After many such tasks, Pepper might unintentionally, without noticing that it's happening, and under guidance from the appropriate statement, develop the ability to infer (in one cognitive step) from a camera's f-stop being lowered to a shallower depth of field in the image. In this case, Pepper adopts the f-stop rule.

[4]One might deny that rules have contents at all. But there are certainly sentences capable of *instructing* us in the following of a rule. Whatever such sentences express is all I mean by the content of a rule.

[5]This is not substantially different from Birman's definition: 'By "adopt" here we mean that the subject...picks up a way of inferring according to [a logical principle], something he wasn't able to do before, *on the basis* of the *acceptance* of the corresponding logical principle' (Birman 2015, p. 31, emphasis in original).

[6]A camera's f-stop is an inverse measure of the width of its aperture and so of its exposure to the light in its environment.

[7]Nor is the issue rationality. Surely, rule acquisition by social immersion is at least sometimes rational.

A rule's acquisition being deliberate or conscious is also not sufficient for its adoption. Consider the case of Oshi, who never picked up the rule of disjunctive syllogism. In order to acquire it, he takes a pill that changes his brain chemistry in just the right way so that he thereafter knows how to use it. This is a case of deliberate and conscious rule acquisition. Nevertheless, it is not a case of adoption, since the acquisition did not involve an explicit representation of the rule. Even if we stipulate that the pill represents the rule to him as it changes his brain chemistry, the case is still not one of adoption. The representation would need to guide him in the process for the acquisition to count as adoption.

Let's move to the crucial matter of the *kind* of guidance offered by these explicit representations (usually sentences). The natural—and, I'll argue, correct—thought is that it is inferential guidance: these sentences guide us along the adoption process by facilitating inferences that move the process forward. This is, after all, how sentences *typically* guide us along epistemic processes: by inferring from them, we move closer to the end of the process.

Consider again the case of Irene, who is learning how the knight moves. Irene has, let's say, a knight on b3 and is wondering whether it is legal to move it to c5. She reminds herself of the rule: 'The knight can move two steps in any direction and then one step in a perpendicular direction, so long as the destination is unoccupied'. She then reasons, 'well, this unoccupied square is two steps up and one step to the right of where my knight is now. So it can move there!'

As the 'so' indicates, the statement of the rule guides Irene by allowing her to make an inference from it. This is not a problematically cherry-picked example. Irene's case is paradigmatic of the kind of process we undergo when using the statement of an unfamiliar rule in order to become competent with the rule itself.

This case is not meant to be conclusive, but only to show that inferential guidance is prima facie the kind of guidance involved in rule adoption. We should accept it as the correct kind of guidance unless there is a problem with doing so. So let's now consider what may at first seem to be a devastating objection to it, raised by Boghossian and Wright (2023).

Regress objection: The fact that adopting R requires inferential guidance by a sentence, S_R, that expresses the content of R, assumes that inferring *in general* requires this kind of guidance. But then, as argued by Boghossian (2014), in order to R_1-infer, one must first R_2-infer from S_{R_1}; but in order to R_2-infer, one must first R_3-infer from S_{R_2}; and so on. So we have a vicious regress much like the one familiar from Carroll (1895).

The right response here is to reject the 'assumes that' claim. Inference *never* requires inferential guidance by a sentence. This already follows from inferences being moves of one cognitive step. Part of what distinguishes one's ability to R-infer from one's mere knowledge of the content of R is whether one needs to make use (inferential or otherwise) of an explicit representation of that content in order to move from R-premises to the R-conclusion. One only R-infers when one does *not* require intermediate use of an explicit representation.

What explains the need for inferential guidance by S_R, then? It is precisely that we wish to arrive at an R-conclusion without using R. If we could use R, we wouldn't need this guidance. Put another way, the regress objection only arises if one conflates adopting R

with using R. If one assumes that the guiding inferences from S_R mutually constitute *a use of* R, it is natural to worry that this can only be motivated if use of R in general requires such guidance. But the guidance is only needed because we *aren't* using R.

But if we don't use R even when we successfully move from R-premises to R-conclusion under guidance from S_R, then in what sense have we adopted it? The answer is that adoption of R is a diachronic *process* comprising numerous cognitive movements from R-premises to R-conclusion. Though these movements all involve inferential guidance by S_R, none of them involve use of R. Rather, the ability to use R is the outcome at the end of the process. It is what we can do after adoption is complete. Let's put this more precisely, as a definition:

Agent W adopts inference rule R during interval $[t_1, t_2]$ iff:

(i) Prior to t_1, W is unable to systematically[8] move from R-premises to corresponding R-conclusions;

(ii) Between t_1 and t_2, W is systematically able to move from R-premises to R-conclusion as a result of inferential *guidance* from S_R, but is unable to use R; and

(iii) After t_2, W is able to use R.

Consider again Irene learning how the knight moves. At the beginning of the process, she is not able to use the knight rule. During the process she is still not able to use it, though she manages to reach its conclusions by inferring—using rules she *is* competent with, like modus ponens—from a statement of it. At some point after doing this enough times, she no longer needs to remind herself of the statement of the rule nor take any intermediate cognitive steps. She develops a new ability to, so to speak, see where she can move her knight. It is only then that she has learned to *use* the knight rule.

The distinction between using R and inferentially moving from R-premises to the appropriate R-conclusion is crucial to this account. One can inferentially move from R-premises to R-conclusion without using R if one uses other rules and a statement of R to bridge the same gap in multiple cognitive steps.[9]

To sum up, in very general terms: adoption is a process of learning a new ability. This process involves using abilities one already has, along with some supplementary guidance (viz. the statement of the rule), to attain the same outcomes as the target ability. Only after so doing many times does one gain the target ability and no longer need the guidance. This kind of learning process is both intuitive and familiar from other learning processes, such as perceptual learning and grammar acquisition.[10] It is, in part, a process of acquiring a disposition—a process that typically involves learning-by-repetition.

We now have a definition of adoption on the table that does away with the regress worry and

[8] I will say more about the 'systematically' qualification in due course.

[9] This is not the *only* way to do so, merely the way that is relevant to my account. One can also inferentially move from R-premises to R-conclusion without using R—even in one step—if this movement is also an instance of another rule, or if one infers according to another rule incorrectly.

[10] See Waite et al. (2019) and Ambridge and Lieven (2011), respectively. The point that rule adoption is a process of skill acquisition by repetition is also made by Devitt and Roberts (2023).

is psychologically plausible. In order to see what it predicts the basic rules are, and what it implies about their adoptability, we need to look more carefully at the semantic structure of sentences that express inference rules.

3 *R*-journeys

Consider the rule of conjunctive syllogism (CS): $\phi, \neg(\phi \wedge \psi) \Rightarrow \neg\psi$. We can express it in English as:

(CS-NL) For any two sentences, if the first is true and the negation of their conjunction is true, then the negation of the second is true.

Let's translate this to a more perspicuous philosopher's English:[11]

(CS-S) For all x, y: If x is true and $\dot\neg(x \dot\wedge y)$ is true, then $\dot\neg y$ is true.[12]

Now, suppose I do not have the rule CS. You want to help me, so you assert (and I accept) a sentence semantically equivalent to (CS-S). To see if I've got it, you also tell me:

(Cat) Gingerface is a cat.

(GF) Gingerface isn't both a cat and a dog.

Given that I cannot use CS, what can I do to get from these three premises to the intended conclusion that Gingerface isn't a dog? The obvious path is as follows. First, use UI (twice) on (CS-S) to infer (using italics as quotation for readability):

(GF Instance) If *Gingerface is a cat* is true and $\dot\neg$(*Gingerface is a cat* $\dot\wedge$ *Gingerface is a dog*) is true, then $\dot\neg$*Gingerface is a dog* is true.

I must now exercise competence with dotted functions. What this involves in natural language is the ability to move from claims like 'The negation of "A" is true' to '"not-A" is true', and its analogues. I'll call this ability *Syntactic competence* (SC). Using it, I can infer:

(GF Conditional) If *Gingerface is a cat* is true and *Gingerface isn't both a dog and a cat* is true, then *Gingerface isn't a dog* is true.

I can then use truth introduction ($\phi \Rightarrow Tr\ulcorner\phi\urcorner$) on (GF) and (Cat), respectively:

(Cat-Tr) *Gingerface is a cat* is true.

(GF-Tr) *Gingerface isn't both a cat and a dog* is true.

Now I can use the rule of Adjunction ($\phi, \psi \Rightarrow \phi \wedge \psi$) on (Cat-Tr) and (GF-Tr) to obtain:

(GF Conjunction) *Gingerface is a cat* is true and *Gingerface isn't both a cat and a dog* is true.

I then use modus ponens (MP) on (GF Conditional) and (GF Conjunction) to infer:

[11]'$\dot\neg$' denotes a function from a sentence to its negation (and analogously for other dotted operators).

[12]I suppress the restriction to sentences for simplicity. This makes no difference, so long as 'a is true' is false when 'a' doesn't denote a sentence; and the dotted functions take non-sentences to non-sentences.

(Dog-Tr) *Gingerface isn't a dog* is true.

Finally, I can use truth elimination ($Tr\ulcorner\phi\urcorner \Rightarrow \phi$) to arrive at the intended conclusion:

(Dog) Gingerface isn't a dog.

Call this kind of inferentially guided movement an *R-journey*: the shortest and most natural inferential path from R-premises and a sentence that expresses R, to the appropriate R-conclusion. While we needn't always take the shortest and most natural inferential path during adoption, it is a harmless idealization to proceed on the assumption that we do.

Note that six rules play a special role: UI, MP, Adjunction, SC, and the truth rules Tr-I and Tr-E. This is not unique to this example, but falls out of the general schema, implicit in (CS-S), for how sentences *in general* express rules. So, these six are the basic rules.

This schema reflects the structure of rules. All rules are general; except in the case of laws, they are conditional; and, except in cases of single-premise rules, their premises behave conjunctively. The obvious way to cash out these features is via quantification (over syntax, when needed); a (perhaps non-material) conditional; and conjunction, respectively.[13]

The list of basic rules needs to be qualified in three important ways, however. Adjunction is only used when R involves multiple premises. Assuming we take a law to be a zero-premise rule, MP is only used when the rule being adopted isn't a law. Finally, SC and Tr-I/E are only used when R is general over sentences or other linguistic elements. Consider, for example, an explicit statement of the knight rule, which might look something like this:[14]

(Knight) For all x, y, z: If x is a knight, y and z are squares two-and-one squares apart, x is at y, and z is unoccupied, then x may move to z.

Evidently, an R-journey from (Knight) and some appropriate premises (e.g. 'this is a knight', 'b3 and c5 are two-and-one squares apart', etc.) requires use of UI, MP, and Adjunction, but not Tr-I/E or SC.[15]

Let's revisit the definition of adoption offered in §2 in light of this account of inferential guidance. The adoption process begins when we accept—be it as a result of testimony or deliberation—a statement of R. Doing so allows us to inferentially traverse from R-premises to R-conclusions, as just modeled above. This is where the content-carrying sentence does its work of inferential guidance. After a number of repeated such paths, we learn to make the

[13]Rules are also normative. It isn't clear whether their normativity, strictly speaking, features in their content. If it does, it can be added to the general schema of rule-expressing sentences. This may require adding another basic rule—in effect, an elimination rule for the normativity. I'm setting this aside, as it doesn't affect the rest of the account and raises difficult tangential issues.

[14]It would have to be a bit more complicated, to account for, e.g., the impossibility of moving a knight that is blocking a check. For simplicity, I ignore these complications.

[15]Cf. Cohnitz and Nicolai (forthcoming), who propose a related but not identical set of basic rules. Notably, they avoid the need for UI and truth by using a schematic substitution principle in place of UI. One important drawback of this is that it does not allow for the adoption of non-linguistic rules like the knight rule. Though there is some interest in investigating what is needed for adopting linguistic rules in particular, this approach does not do justice to the ubiquitousness of adoption that was defended in §2 and which is crucial to the 'indispensability argument' that will be presented in §4.1.

jump from R-premises to R-conclusion in one shot. It is when we so learn that the adoption process is complete. We can now *use* R.

It is not a problem for my account if, even after the learning process is over, we still require non-inferential guidance by a (non-explicit) representation of the rule in order to use it.[16] So long as we no longer need inferential guidance by an explicit representation, and so can proceed in one cognitive step, I can be neutral about the mechanics of that cognitive step.

One may ask: if non-inferential guidance is still required for inference, why posit the need for inferential guidance in addition? The answer is that inferential guidance isn't an additional requirement on inference; it's a requirement on *adoption*. There's nothing surprising about distinct activities requiring different kinds of guidance.

Another potential worry is a version of the rule-following problem familiar from Kripke (1982). Learning to use R as a result of finitely many moves from R-premises to R-conclusion must involve some kind of recognition (perhaps subpersonal, perhaps not) that these moves all conform to a formal pattern: the pattern sanctioned by R. But these finitely many moves conform to infinitely many formal patterns. How do our cognitive processes pick out the right one?

My account provides a highly intuitive answer to that question: it is the one expressed by the sentence we've accepted and which has guided these inferential movements. In other words, the rule-expressing sentence plays a double role. It inferentially guides us from R-premises to R-conclusion *and* it identifies the rule, of the infinitely many options, that constitutes a short-cut through the finitely many R-journeys made during the adoption process.

4 The Kripke-Birman argument

4.1 Unadoptability and indispensability

We now have a list of basic rules: UI, MP, Adjunction, SC, and Tr-I/E. It is tempting, given the R-journey model of adoption, to accept the following:

> **(Basic Rule Indispensability) (BRI)** One is able to infer from R-premises to the corresponding R-conclusion based on inferential guidance from a statement of R only if one is antecedently able to use the six basic rules (or equivalents).[17]

The qualification 'or equivalents' is meant to stave off a concern about someone who, say, is not able to use MP but is able to use rules that, together, mutually entail every instance of MP (e.g. where the conditional is material, MP is derivable from explosion ($\phi, \neg\phi \Rightarrow \psi$) and proof by cases).

With (BRI) to hand, one can run arguments for the unadoptability of basic rules. Here's the argument for UI:

[16] As argued by, e.g., Boghossian (2014).

[17] Officially, this needs to come with the caveat that Adjunction is only needed for rules with multiple premises, MP needed only for rules that aren't laws, and SC and Tr-I/E only for linguistic rules.

Unadoptability argument: Suppose Irene adopts UI during $[t_1, t_2]$. Then she is unable to use UI prior to t_2 (from (i) and (ii) in the definition of adoption). She is also systematically able to infer from UI premises to UI conclusions during $[t_1, t_2]$ under inferential guidance from a statement of UI (from (ii)). But she can only do this if she is antecedently able to use UI (from (BRI)). Contradiction. So Irene does not adopt UI during $[t_1, t_2]$.

Since MP isn't a law, Adjunction involves two premises, and SC and Tr-I/E are linguistic rules, the unadoptability argument works for them as well. I believe this is a charitable and indeed quite plausible way to interpret the Kripke-Birman argument.

Note that the key step in the argument is the use of (BRI), which expresses the indispensability of basic rules for adoption of rules *in general*. It is the indispensability that does the crucial work. This point can be exploited to mount an argument against *dropping* basic rules. Unlike the unadoptability argument, which targets the *possibility* of adding basic rules, the argument against dropping basic rules is normative:[18]

Indispensability argument: Adoption of rules in general requires the basic rules. We must not lose the ability to adopt new rules, on pain of severe cognitive impairment. So we must not drop any of the basic rules.

This argument has significantly more bite than the unadoptability argument. Most if not all rational agents already have the ability to use the basic rules, so the unadoptability argument only applies to merely possible cases. The indispensability argument, on the other hand, targets plenty of live proposals. The free logic response to puzzles of non-existence (e.g. Sainsbury (2005)) drops unrestricted UI; some solutions to the semantic and set theoretic paradoxes (e.g. Priest (2017)) drop unrestricted MP, as does McGee's solution to his famous election case (McGee (1985)); the subvaluationist solution to the sorites paradox (Hyde (1997)) drops Adjunction; and *all* classical approaches to the liar paradox drop Tr-I/E.

Now, one may worry that I am misinterpreting Kripke's argument. And, in one sense, I obviously am. Kripke is clear that we should not take adoption considerations to show that certain specific rules are unadoptable:

> But, it will be said, maybe this doesn't apply to all of logic. Maybe there is some more basic ur-logic for which these kinds of arguments [don't] apply.... Such a reply doesn't really get the point.... One has to just think, not in terms of some formal set of postulates but intuitively, that is, one has to *reason*. One can't just 'adopt' a formal system independently of any reasoning about it because if one tried to do so, one wouldn't understand the directions for setting up the system itself. (Kripke 2023, p. 20, emphasis in original)

One way to understand Kripke here is as reminding us that reasoning happens in our metalanguage. That we can pull some trick with a defined object language tells us little about what we can do with our *actual reasoning*. Though I am sympathetic to this line of thought, it does not have the implications Kripke seems to want it to have. To be sure, one does need to reason in order to follow 'the directions for setting up' a system of inference rules.

[18]I haven't offered a model for how rules can be dropped, so one may wonder whether dropping rules is possible. I'll address this in §7.2.

But nothing stops us from inquiring into exactly what kind of reasoning is needed to do just that. That is, we can, after reasoning, step back and study that reasoning. The account of adoption I've offered, along with the Kripke-Birman argument as I've reconstructed it, are steps in that direction. And they do point to a special place for the basic rules in the process of reasoning with such directions.

Much more can be said here. What's clear is that Kripke (2023) offers numerous arguments and considerations against the possibility of adopting a logic. My reconstruction is meant to capture *one type* of argument he gives: the raven argument (ibid., p. 15–16), which was notably reconstructed by Berger (2011, p. 183–184), adapted by Birman (2015) into the Harry argument, and which has been the subject of much recent debate, as in: Finn (2019), Cohnitz and Nicolai (forthcoming), Birman (2023), and Boghossian and Wright (2023). *This* argument type surely targets specific rules: precisely those used in R-journeys.

As a final exegetical point, Kripke himself (ibid., p. 18) identifies MP and Adjunction as the other rules besides UI to which his raven argument applies.[19] And in fn. 18, he notes the need for 'Tarskian disquotation' as well.

4.2 Why the arguments fail

The problem with both the indispensability and unadoptability arguments is that (BRI), even with the caveat about 'equivalents', is false. It is just not the case that UI, MP, Adjunction, SC, and Tr-I/E are indispensable to the process laid out in §3. This is because restrictions of these rules suffice for the process to carry through. Nor, as we'll see, is this objection avoidable by targeting the restrictions instead.

Let's consider the case of MP.[20] Suppose Irene is only competent with MP restricted to conditionals that don't contain embedded conditionals. This means Irene is unable to move from premises of the form ϕ and $\phi \to (\psi \to \chi)$ to the conclusion $\psi \to \chi$ in one cognitive step. This does not mean that Irene rejects these MP instances—she is simply not disposed to make them. Indeed, Irene may be able to *derive* some specific instances of this pattern, provided she has additional premises that bridge the inferential gap. Let's assume, however, that if there are any such specific instances, they are exceptions. Prior to adoption, Irene does not have the resources to systematically derive all of her missing instances of MP.

Irene now wants to adopt the full MP. She begins the process by accepting a sentence that expresses unrestricted MP, namely:

> **(MP-S)** For all x, y: If x is true and $x \dot{\to} y$ is true, then y is true.

Suppose Irene accepts arbitrary sentences A and $A \to (B \to C)$. Now that she's accepted (MP-S), Irene can derive $B \to C$, by going through an R-journey of the kind modeled in §3. In particular, she can use Tr-I and Adjunction to infer:

[19] Kripke later makes an argument concerning the law of non-contradiction. But that argument is rather different from the raven-type arguments. This is unsurprising, since there is no obvious place to use that law during an R-journey.

[20] Considerations for other rules are similar, though I'll also discuss the UI case in the next section.

$$\ulcorner A \urcorner \text{ is true and } \ulcorner A \to (B \to C) \urcorner \text{ is true.}$$

From (MP-S), she can then infer (using UI and SC):

$$\text{If } \ulcorner A \urcorner \text{ is true and } \ulcorner A \to (B \to C) \urcorner \text{ is true, then } \ulcorner B \to C \urcorner \text{ is true.}$$

Since this does not syntactically embed a conditional, her restriction of MP suffices to infer its consequent.[21] Tr-E completes the R-journey to $B \to C$. Adoption is complete when Irene takes enough of these R-journeys (using different specific sentences) to, so to speak, pick up the pattern, so that she no longer needs to make intervening derivations.

Is there a restriction of MP sufficient to carry the adoption process forward but which can't itself be restricted without thereby impeding R-journeys? There isn't. To see why, suppose Irene does not have full MP but has the restriction: $\phi \wedge (\phi \to \psi), (\phi \wedge (\phi \to \psi)) \to \psi \Rightarrow \psi$. Now suppose Irene accepts arbitrary sentences $A, A \to B$ that do not correspond to this form. Using adjunction, Irene can infer $A \wedge (A \to B)$. By instantiating on (MP-S) with $\ulcorner A \urcorner$ for x and $\ulcorner B \urcorner$ for y (and then using SC and Tr-E), Irene can infer $(A \wedge (A \to B)) \to B$. Her restriction allows her to infer B. So this restriction suffices to recover unrestricted MP.

This move can be iterated. Letting π abbreviate the form $\phi \wedge (\phi \to \psi)$, the restriction just considered is of the form $\pi, \pi \to \psi \Rightarrow \psi$. We can suppose that Irene does not have this either, but has $\pi \wedge (\pi \to \psi), (\pi \wedge (\pi \to \psi)) \to \psi \Rightarrow \psi$. Evidently, the same argument carries through and can be indefinitely iterated. So any restriction of MP in this series could be further restricted without impeding the possibility of adopting unrestricted MP. Furthermore, we could have started the series with any restriction of MP sufficient for adoption rather than with the full MP. It follows that a restriction of MP which suffices for adoption *but which cannot itself be restricted without impeding adoption* does not exist.

There is, in any case, a more general problem with the purported indispensability of any rule: the adoption process requires only a finite (and usually relatively small) number of R-journeys in order to succeed. Since natural languages contain infinitely many sentences, there is no general form of MP restriction that is indispensable to adoption. Any finite set of MP instances, *so long as it is numerous and varied enough*, will do. The same applies to the other basic rules.

This *numerous and varied enough* condition is what the 'systematically' conveys in (i) and (ii) in the definition of adoption. The idea is that adoption is a process of learning to cognitively shortcut past R-journeys. But if our attempts at traversing R-journeys is interrupted enough by cases in which exceptions to basic rules block the inferential path, the cognitive shortcutting will not occur. So, though we need only relatively few instances of the rules, the batch of available instances must be reasonably robust. It does not follow from this, however, that they can have no exceptions at all—nor that there is a general form of restriction beyond which they cannot be restricted.[22]

[21] If Irene has the Importation rule, viz. $\phi \to (\psi \to \chi) \Rightarrow (\phi \wedge \psi) \to \chi$; and can, as the example assumes, use MP in cases of conjunctive antecedents, then she can always derive MP for embedded conditionals. But it does no harm to the point of the example to suppose that she lacks Importation as well.

[22] Though Cohnitz and Nicolai (forthcoming) challenge the scope of the adoption problem, they concede that at least MP and the schematic substitution principle they employ in place of UI are needed for adoption

It is true that one needs *some* (finite) restrictions of the basic rules in order to successfully and reliably adopt new rules. So an indispensability argument against dropping any of these *in toto* can be mounted (as can an unadoptability argument applied to those who lack any restriction whatsoever). But no one advocates that we drop any of these rules altogether. All serious proposals for alternative logic involve relatively minor restrictions.

It's also true that finite restrictions of basic rules are highly unmotivated. The point is just that, as far as what adoption requires, there is no minimum *general form* of restriction.

This way out of the unadoptability and indispensability arguments rests on my account of adoption as a finite diachronic process. One could try to rescue the arguments by rejecting this picture of adoption, but the prospects for doing so are dim. Inferential guidance by rule-expressing sentences is crucial to these arguments (both in my reconstructions and in Kripke (2023) and Birman (2015)). The role of this guidance is either limited to a finite adoption process, or else is involved in inference with the adopted rule indefinitely. But the second horn is completely unmotivated unless one is committed to the inferential model of ruling following that walks into the Carroll-type regress problem.

There's another problem with trying to rescue the Kripke-Birman argument. Without an account of adoption on which restrictions of basic rules suffice for the adoption process, we need to adopt a subclassical logic to accommodate Tr-I/E in their full generality—or else lose the ability to adopt new linguistic rules altogether. This is a conclusion I assume most fans of the adoption problem would wish to avoid.

5 Anti-exceptionalism

Debates about the adoption problem have sometimes been carried out in the context of debating *anti-exceptionalism about logic*: the view that truths of logic are, at least in principle, open to revision on abductive grounds. The unadoptability argument, it has generally been supposed, poses a problem for anti-exceptionalism. If basic rules of logic can't be adopted *at all*, then they certainly can't be adopted on abductive grounds.[23]

There's an easy reply on behalf of anti-exceptionalism. It is not part of the anti-exceptionalist view that one can rationally theorize without the use of certain basic concepts, like ALL, or some basic forms of reasoning, like (some sufficiently robust restriction of) UI. The charitable reading of the position is much weaker, though still strong enough to be of interest: an agent, assumed to have sufficiently robust reasoning abilities to allow for theorizing in the first place, may and can change any of her beliefs (and corresponding reasoning patterns), including any logical ones, in response to abductive evidence. Cases like that of Harry, who lacks even a restriction of UI, are simply not relevant to a reasonable version of anti-exceptionalism.

in general. They defend this by raising doubts as to whether a logic lacking these rules 'could be an adequate formal model for any possible form of natural reasoning' (p. 14). If my model of adoption is correct, we needn't make such a strong claim about what counts as natural reasoning, since neither of these rules (nor any restrictions of them) are in fact indispensable to adoption.

[23]See Boghossian and Wright (2023, §5). NB: Sometimes anti-exceptionalism is characterized as the view that truths of logic are open to abductive revision based on *empirical* evidence. But whether the abductive considerations are empirical or a priori makes no difference to the general issues at hand.

It is rather the indispensability argument that is relevant to anti-exceptionalism. For, if dropping a basic rule leaves a theorizer so cognitively impaired that she is unable to carry on theorizing, this would, perhaps, trump any abductive considerations in favor of doing so. Relevant though it may be, however, we've already seen that the indispensability argument is not sound. One can drop a rule *in its full generality* without losing all of its instances—and this is what is called for in all actual cases of proposed logical revision. As I've shown, doing so is not cognitively impairing, at least not in the sense of impeding the general ability to adopt new rules.

So, adoption considerations alone do not tell against the potential viability of dropping certain rules of logic, even basic ones, in response to abductive considerations. My account nevertheless affirms a special role for basic rules. UI, MP, Adjunction, SC, and Tr-I/E (or whatever restrictions of them a given agent in fact uses) do play a special role in the adoption process—one that ultimately stems from the special role that ALL, IF, AND, names for syntax, and TRUE play in the general contents of inference rules.

It may help to run through how a case that is relevant to an actual debate in logic might play out. Suppose Irene reasons with a restriction of UI that does not apply to empty names. This could be a curious accident of Irene's language-acquisition process or it could be because Irene intentionally, in response to abductive considerations, restricted UI. Suppose also that Irene is now convinced, based on a (re-)evaluation of the abductive evidence, that she ought to adopt the full UI. That is, she accepts:[24]

> **(UI-S)** For all x, y, t: If $\dot{\forall} yx$ is true, then $subs(x, y, t)$ is true.

Irene, let's say, accepts a specific universal sentence: $\forall v\phi$. She instantiates into (UI-S) to infer (with the required assistance from SC) the following, where 'a' is an empty name:

> **(Empty Instance)** If $\ulcorner \forall v\phi \urcorner$ is true, then $\ulcorner \phi(a/v) \urcorner$ is true.

She is perfectly able to do this, since *the name of* 'a' is not empty. In other words, the names of elements of Irene's syntax—including names for her empty names—are not among the empty names. Tr-I, MP, and Tr-E complete the R-journey in the familiar way. Doing this many times, Irene is able to adopt back full UI through the process modeled above.

The point of the example is this: Irene can restrict UI without losing the ability to adopt new rules, *including adopting back the full UI.* So her restriction of UI does not introduce any special problems that could be used to undermine anti-exceptionalism.

6 When might adoption fail?

I've shown that the basic rules are not indispensable to adoption, since some restrictions of them (including some theoretically motivated ones) suffice for the adoption of new rules.

[24] $\dot{\forall}$ is a function from a variable and a formula to the corresponding universal formula; *subs* is a function from a formula, a variable, and a term, to the formula with the appropriate substitution. As usual, I suppress the implicit antecedent in (UI-S): x is a formula, y a variable, and t a term.

Nor is any such restriction itself indispensable, since there is always a way of further restricting without impeding the possibility of adoption. None of this implies, however, that *all* restrictions of basic rules suffice for adoption. We need only consider the trivial case of the *total* restriction. So, when does a restriction suffice and when doesn't it?[25]

In part, this is a matter of luck. To see why, consider again the case where Irene has UI restricted to non-empty cases. Suppose she is trying to adopt an inference rule that, like the knight and f-stop rules, isn't linguistic. For simplicity, say it is of the form $\phi(t) \Rightarrow \psi(t)$, so that its content-expressing sentence is:

> **(Simple Rule)** For all x: If $\phi(x)$, then $\psi(x)$.

Irene begins with the premise $\phi(a)$, for some specific name a. She must then instantiate on (Simple Rule) with a. If a happens to be an empty name, however, she will be unable to do so. In such a case, Irene is simply unlucky. If she'd picked a different instance of the rule—one not involving an empty name—the R-journey would have gone through fine.

Recall that adoption requires that we run through multiple R-journeys, until the pattern sticks. If one attempted R-journey fails, Irene can keep trying. So, for Irene's attempt at adopting the target rule to fail, she would have to be very unlucky:[26] she would have to repeatedly attempt R-journeys using exactly the instances that her restriction of UI blocks. This could happen—it would be unreasonable to expect attempts at adoption to be infallible. But it requires unusual circumstances.

Are there failure cases that aren't a matter of luck, but are systematic? This would obtain if all instances of a basic rule used in adoption belonged to a privileged *core restriction* of the rule. Then a restriction of a basic rule that eliminated all instances of that rule's core restriction would not suffice for adoption. However, there aren't any core restrictions, since inference rules can govern arbitrary kinds of situations and vocabulary. For example, a restriction of UI that eliminates all instances except those to names of knights and of squares on chess boards suffices for adopting the knight rule.

Even so, there is some interest in identifying what is needed to adopt arbitrary *linguistic* inference rules. For this class, the following core restrictions can be identified: UI restricted to names of syntax; MP restricted to cases in which the consequent is a truth ascription and the antecedent is either a truth ascription or a conjunction of truth ascriptions; and Adjunction restricted to truth ascriptions.[27] All instances of a basic rule used during the adoption of a linguistic rule belong to the core restriction of that basic rule.

Evidently, a restriction of any basic rule that lacks *all* instances of its core restriction is insufficient for the adoption of linguistic inference rules. This is unsurprising, since losing the core restriction in full amounts to an almost total elimination of the ability to apply the rule in question to reasoning about syntax. This does not mean that we can't, without impeding adoption, restrict a core restriction of a basic rule *at all.* The arguments from §4.2, to the effect that restrictions sufficient for adoption can always be further restricted, apply

[25]Thanks to an anonymous referee for raising a number of the issues addressed in this section.

[26]Or sabotaged, or masochistic, etc.

[27]The truth rules and SC, which are both only needed for linguistic rules, are their own core restrictions.

just as well to core restrictions. In particular, since adoption requires only finitely many R-journeys, only (sufficiently numerous and varied) finite subsets of the core restrictions are needed to successfully carry out the adoption of linguistic rules.

So far, we've established that fully eliminating a core restriction of a basic rule blocks adoption, though there are (infinitely many) restrictions of the core restrictions that do not. This leaves open: are there theoretically motivated restrictions that cut out enough of the core restriction to systematically impede adoption? Let's consider a candidate.

Suppose Irene adopts the logic LP in response to the semantic paradoxes, thereby restricting MP to non-paradoxical instances. Most of those who advocate LP propose that it be augmented with a non-truth-functional conditional that satisfies unrestricted MP. For the sake of the case, however, assume Irene does not do this; or that, if she does, she nevertheless expresses the contents of rules via the material conditional. Finally, assume that, as is typical for non-classical approaches to truth, truth ascriptions of paradoxical sentences (and their conjunctions) are themselves paradoxical.[28] It's easy to see that if Irene tries to get back unrestricted MP by running through R-journeys using (MP-S) and paradoxical premises, she will be blocked at the MP step.

This case suggests a general condition for when adoption might be systematically blocked. Suppose an agent has a restriction, B', of basic rule B, and the agent is trying to adopt rule R. Let Γ_R be a set of premises corresponding to an instance of R. R-journeys involve certain successive operations on the starting premises plus R-S (the statement of R): instantiation, truth introduction, adjunction, etc. Let $\Phi_B(\Gamma_R)$ be the set containing the result of those operations *that have to occur prior to the use of basic rule B in R-journeys* on the members of $\Gamma_R \cup \{\text{R-S}\}$. Adoption *might* be systematically blocked only in cases where for all Γ_R corresponding to instances of R that the agent is missing, the members of $\Phi_B(\Gamma_R)$ needed for the B-step of the R-journey correspond to instances of B missing from B'. In the Irene case, R and B are both unrestricted MP, and B' is the restriction of MP to non-paradoxical instances. Adoption might fail because the instance of MP we get when we instantiate (MP-S) with names of paradoxical MP-premises and perform the needed operations on the starting premises is, itself, an instance of MP with paradoxical premises.

I have stressed that even when this condition holds, adoption only *might* be blocked. To see why, consider that Irene can still go through R-journeys using non-paradoxical premises. In these R-journeys, she would use (MP-S) to move from MP premises to conclusion in multiple steps despite having the ability to do so in one step (using her restricted MP). Even so, these MP-journeys needn't be idle, since (as argued at the end of §3) (MP-S) plays a dual role: besides providing inferential guidance, it also fixes the rule, out of the infinitely many candidates, that corresponds to the pattern of the MP-journeys one has taken. Since (MP-S) expresses unrestricted MP, this is the rule that is selected.

Whether the suggestion of the preceding paragraph works requires empirical investigation. The general model of adoption on offer assumes that by repeatedly moving through an inferential pattern, under guidance by a statement that expresses that pattern, one learns to infer according to that pattern. As argued in §3, this is both intuitively plausible and

[28] Priest (2017) non-committally explores an account in the neighborhood of these conditions.

empirically supported. But we are now up against a narrower empirical question, which we can put as follows. Consider an agent who is competent in a certain inferential pattern, R', which is a restriction of a more general pattern, R. Suppose the agent repeats the pattern R' under inferential guidance from a sentence that expresses R: does this suffice for the agent to become competent in the full R? The answer to this empirical question is not at all obvious; nor is it obvious that it even has a general answer. It may well depend on the nature of R and R', or on other factors.

7 Boghossian and Wright

Aside from the regress problem for the inferential model of rule-following, Boghossian and Wright (2023) raise a number of issues having to do with the adoption problem. I do not have the space to discuss most of them, but there are two that could be seen as problems for my proposal and so which I would like to briefly touch on.

7.1 The lambda operator adoption problem

Though Boghossian and Wright (hereafter B&W) reject the Kripke-Birman argument, primarily as a result of the regress problem, they believe a different argument can be mounted against the possibility of adopting *any* of the standard rules governing the logical connectives and quantifiers. They call it the *lambda operator adoption problem*, and argue that it works independently of the kind of guidance that statements of rules are meant to provide during adoption.

Call an operator, λ, a *lambda operator* just if 'a canonical statement of its characteristic inferential profile must essentially involve (metatheoretic) use of (a counterpart of) λ itself' (Boghossian and Wright 2023, p. 100).[29] B&W's strategy is to replace the assumption of inferential guidance in the Kripke-Birman argument with the assumption that the rule being adopted, R, is a rule governing a lambda operator, λ. The thought is then that a canonical statement of R can only guide the adopting agent toward inferential behavior that is implicit in her understanding of the instance of λ that occurs in the statement of R. But if it is so

[29]B&W argue that the standard logical operators and quantifiers are all lambda operators—or, at least, those chosen as 'semantical primitives' are (ibid., p. 101). But this is just not so. The 'characteristic inferential profile' of 'not', for example, is given by (the conjunction of) the following (using schemata in place of quantification over syntax for simplicity):

($\neg$-I-S) If ϕ entails an absurdity, then $\neg\phi$ is true.

($\neg$-E-S) If ϕ is true and $\neg\phi$ is true, then ψ is true.

(DNE-S) If $\neg\neg\phi$ is true, then ϕ is true.

The conjunction of these statements does not involve a metatheoretic use of 'not'. One might try to argue that $\neg$-E is better expressed as 'It is not the case that: ϕ is true and $\neg\phi$ is true'. But if this were the statement of $\neg$-E, we could not distinguish between Explosion (i.e. $\neg$-E) and the Law of Non-contradiction (i.e. $\neg(\phi \wedge \neg\phi)$). Many who reject the former (e.g. dialetheists and relevance logicians) accept the latter.

Unsurprisingly, a rule-by-rule inspection of canonical statements of standard rules for logical operators reveals there to be only three logical lambda operators: 'all', 'if', and 'and'. This has nothing to do with choice of primitives, and everything to do with their role in the general contents of inference rules.

implicit, then adoption is not needed; at best, the agent is drawing out a rule she's already committed to, not adopting a new one. So adoption of a rule governing a lambda operator is not possible.[30]

Since this argument is meant to apply to adoption via *any* kind of guidance by content-expressing sentences—not limited to, but including inferential guidance—the model I've offered of R-journeys is enough to show that the argument is unsound. The problematic assumption is that if a canonical statement is to guide one's change in inferential practice, then the use made of λ in the canonical statement must already reflect that change. But there is no reason to think this. An agent can use whatever rules she already has for λ to, in conjunction with the rest of the canonical statement, draw out the conclusion of the hitherto unavailable inferences.

Consideration of quotidian cases make this especially clear. Consider the rule, briefly discussed in §2 (with generality omitted for simplicity):

> **(F-stop Rule)** If an f-stop is lowered, the depth of field becomes shallower.

(F-stop Rule) expresses what for at least some is an inference rule, as discussed in §2. Since its canonical statement uses the predicate 'f-stop', the latter, though not an operator, has the property of lambda operators that does the crucial work in B&W's argument. If the latter were sound, then this statement could only help one reach the conclusion that a depth of field is shallower, from the premise that an f-stop is lower, if the validity of the f-stop rule were already implicit in one's understanding of the 'f-stop' that appears in the statement. But this is obviously not so: MP bridges the gap.

This point applies just as well to logical operators, even those that are genuinely lambda operators, viz. *all*, *if*, and *and*. For example, one can move from the statement of unrestricted UI (viz. (UI-S)) and a universal claim, to an instance, *even if* the 'all' that is used in (UI-S) is restricted so as not to license that particular instance. This was already shown in §5 above, and for MP in §4.2. In general: the rest of the statement of the rule, in conjunction with other rules one is competent with, fills in for what the lambda operator cannot yet do.

Note that B&W's argument would be quite plausible if a canonical statement had to guide one to the right inferential practice without any intervening cognitive work, particularly of an inferential sort. But as we've already seen, adoption is a learning process and, as such, should be expected to involve relatively complex inferential work.

[30]'[I]f all one gets by way of explanation of the new inferential practices [R] ordains is just a statement of R, one has no basis but to interpret that statement in terms of the very notion of λ that it is supposed to supersede, and so will be offered no path to the novel practices the rule is supposed to encode' (ibid., p. 100). And: 'any...canonical statement which is to non-conservatively extend the relevant agent's prior inferential practice with λ must, if it is to codify and explain such a change, draw on an understanding of the use made of λ in its very formulation which *already* diverges from that possessed by its recipients and is thus not available to them at the crucial point' (ibid., p. 100–101, emphasis in original).

7.2 Adopting a restriction

I've made much of the fact that most live proposals for logical revision are for *restricting* a rule. But I've given no account of how one goes about doing so. B&W argue that, at least in the case of UI and MP, restriction is impossible. I will focus on the case of UI, though the argument is analogous for MP. Consider again the statement of UI:

> **(UI-S)** For all x, y, t: If $\dot{\forall} yx$ is true, then $subs(x, y, t)$ is true.

And suppose that we wish to restrict it so that instances into empty names are blocked:[31]

> **(UI*-S)** For all x, y, t: If $\dot{\forall} yx$ is true and t isn't an empty name, then $subs(x, y, t)$ is true.

B&W's point is that (UI*-S) features—supposing one hasn't yet made the restriction—an unrestricted 'all'. (UI*-S) therefore cannot guide one to restrict UI, since that guidance would involve unrestricted UI.

This argument is plausible if we assume that accepting (UI*-S) constitutes, or at least suffices for, restricting UI. But, as with adopting one, restricting a rule is a learning process of which accepting the sentence expressing the restriction is only the first step.

I'll now offer a model of rule restriction. As with my model of rule adoption, it is an idealized model, though one meant to be psychologically plausible.

When we adopt a restriction—or, for that matter, drop a rule altogether—what we do is we train ourselves to stop inferring in a way we're disposed to do. Rather than gain a disposition, we are trying to lose one. We do this by checking our inferences when we make them (in relevant contexts) to see if they are of the kind we are trying to avoid. When we find one that is, we reject it.

Suppose Irene has recently accepted a paracomplete response to the liar paradox. She is now working, perhaps for one of the first times, through a Curry paradox and reaches the point in the derivation where she infers that the Curry sentence, κ, entails $\perp$. So, she reasons: $\neg\kappa$. Knowing that she's still learning to reason paracompletely, she asks herself: was that an instance of a paracompletely invalid rule? Upon noticing that it was, she *rejects* the inference (and its conclusion). Note that she here uses unrestricted $\neg$-I and then rejects that use. She cannot abstain from using the unrestricted rule altogether, since she is still disposed to do so. Only after going through enough reasoning journeys of this sort, she successfully learns to stop reasoning according to unrestricted $\neg$-I when in paradoxical contexts.

Given this model, let's consider the case of UI. Is there an obstacle to restricting it to, say, non-empty instances? It's not clear why there should be, since Irene needn't use the sentence (UI*-S) during the restriction process. What she does instead is simply use UI—a rule she already has—and then reject that use after the fact if it isn't also an instance of the restriction.

[31] B&W do not use this, or any, specific example. But the example conforms to the general schema of their argument. See Boghossian and Wright (2023, p. 105–108).

Now, on my model, Irene does implicitly infer from a sentence like the following, which features the (unrestricted) 'all' she already has:

(UI-Rej) For all inferences x: If x is a use of UI into an empty name, reject x.

But this needn't impede the restriction process. Repeated infer-then-reject journeys are helping Irene learn to stop using empty instances of UI. That she needs to use UI along the way does not undermine the process unless the instances of UI she must use along the way are themselves of the sort she is trying to stop using. But they are not: names of inferences are not empty names.

In sum: the process of restricting a rule one is disposed to use unrestrictedly is, I propose, one of repeatedly—when in relevant contexts—inferring, checking, and then, when appropriate, rejecting. We should expect adoption-type problems to arise if this process itself requires using instances of the rule that go beyond the target restriction. Perhaps some proposed restrictions of basic rules are like this, in which case those proposals do face a potential problem. But there is no reason to think all or most theoretically motivated restrictions of UI or MP are like this.

References

Ambridge, Ben and Elena V. M. Lieven (2011). 'Theoretical Approaches to Grammar Acquisition'. In: *Child Language Acquisition: Contrasting Theoretical Approaches*. Cambridge University Press, pp. 103–136. DOI: 10.1017/CBO9780511975073.005.

Berger, Alan (2011). 'Kripke on the Incoherency of Adopting a Logic'. In: *Saul Kripke*. Ed. by Alan Berger. Cambridge University Press, pp. 177–207.

Birman, Romina (2015). *What the Tortoise Said to Kripke: the Adoption Problem and the Epistemology of Logic*. CUNY Academic Works.

– (2023). 'The Adoption Problem and the Epistemology of Logic'. In: *Mind* 529, pp. 37–60. DOI: 10.1093/mind/fzad009.

Boghossian, Paul (2014). 'What is Inference?' In: *Philosophical Studies* 169.1, pp. 1–18. DOI: 10.1007/s11098-012-9903-x.

Boghossian, Paul and Crispin Wright (2023). 'Kripke, Quine, the "Adoption Problem" and the Empirical Conception of Logic'. In: *Mind*. DOI: 10.1093/mind/fzad008.

Carroll, Lewis (1895). 'What the Tortoise Said to Achilles'. In: *Mind* 4.14, pp. 278–280. DOI: 10.1093/mind/IV.14.278.

Cohnitz, Daniel and Carlo Nicolai (forthcoming). 'How to Adopt a Logic'. In: *Dialectica*.

Devitt, Michael and Jillian Rose Roberts (2023). 'Changing Our Logic: A Quinean Perspective'. In: *Mind*. DOI: 10.1093/mind/fzad010.

Finn, Suki (2019). 'Limiting Logical Pluralism'. In: *Synthese* 198.Suppl 20, pp. 4905–4923. DOI: 10.1007/s11229-019-02134-8.

Hyde, Dominic (1997). 'From Heaps and Gaps to Heaps of Gluts'. In: *Mind* 106.424, pp. 641–660. DOI: 10.1093/mind/106.424.641.

Kripke, Saul A. (1982). *Wittgenstein on Rules and Private Language: An Elementary Exposition*. Harvard University Press.

– (2023). 'The Question of Logic'. In: *Mind*. DOI: 10.1093/mind/fzad008.

McGee, Vann (1985). 'A Counterexample to Modus Ponens'. In: *Journal of Philosophy* 82.9, pp. 462–471.

Priest, Graham (2017). 'What If? The Exploration of an Idea'. In: *Australasian Journal of Logic* 14.1. DOI: 10.26686/ajl.v14i1.4028.

Sainsbury, Mark (2005). *Reference Without Referents*. Oxford, England and New York, NY, USA: Clarendon Press.

Waite, Stephen et al. (2019). 'Analysis of Perceptual Expertise in Radiology – Current Knowledge and a New Perspective'. In: *Frontiers in Human Neuroscience* 13. DOI: 10.3389/fnhum.2019.00213.