Kripke, Quine, the 'Adoption Problem' and the Empirical Conception of Logic

Paul Boghossian
New York University, United States
paul.boghossian@nyu.edu

Crispin Wright
New York University, United States; University of Stirling, United Kingdom. cjgwright@gmail.com

Recently, there has been a significant upsurge of interest in what has come to be known as the 'Adoption Problem', first developed by Saul Kripke in 1974. The problem purports to raise a difficulty for Quine's anti-exceptionalist conception of logic. In what follows, we first offer a statement of the problem and argue that, so understood, it depends upon natural but resistible assumptions. We then use that discussion as a springboard for developing a different adoption problem, arguing that, for a significant class of basic logical principles, there is indeed a difficulty in seeing how they might be 'freely adopted,' thereby vindicating something close to the spirit of Kripke's original claim. This first part of our argument will enforce a significant qualification of Quine's claim that basic logical principles can be empirically confirmed. In the concluding sections of the paper, we turn to the question, specifically, of the empirical revisability of logic, arguing that when proper attention is paid to the role of reasoning in theory revision, it does indeed emerge that anti-exceptionalism, in full generality, is untenable.

1. Introduction

Intuitively, there is a robust distinction between those beliefs, or inferences, that can be justified *a priori* and those that can only be justified *a posteriori*. Sensory evidence is needed to justifiably believe

(1) It's sunny

but not, it would appear,

(2) If it's sunny then it's sunny.

The latter proposition seems, rather, to be 'epistemically analytic', justifiably believable merely in virtue of being understood (see Boghossian 2017 and Boghossian chapters 7 and 9 in Boghossian and Williamson

Mind, Vol. 133 . 529 . January 2024
https://doi.org/10.1093/mind/fzad011

2020). Similarly, sensory evidence would be needed to justify the inference

(4) It's sunny/So, It's warm

but not the inference

(5) It's sunny/So, It's sunny.

The latter inference seems, once again, epistemically analytic.

It is useful to distinguish between *weak a priori* justification and *strong a priori* justification (see Field 2005). If the belief that P has weak *a priori* justification that means that it has been justified without the benefit of sensory evidence;[1] if it has strong *a priori* justification, that means it has weak *a priori* justification and, furthermore, cannot be refuted by purely empirical considerations.

In *Two Dogmas of Empiricism* (1951), Quine outlined a conception of the epistemology of logic according to which no logical truth or inference can be either strongly or weakly *a priori* justified. It is memorably summarized in the following famous passage:

> Any statement can be held true come what may, if we make drastic enough adjustments elsewhere in the system. Even a statement very close to the periphery can be held true in the face of recalcitrant experience by pleading hallucination or by amending certain statements of the kind called logical laws. Conversely, by the same token, no statement is immune to revision. Revision even of the logical law of the excluded middle has been proposed as a means of simplifying quantum mechanics; and what difference is there in principle between such a shift and the shift whereby Kepler superseded Ptolemy, or Einstein Newton, or Darwin Aristotle?[2]

In this passage, the rejection of strong *a priority* is explicit; and the rejection even of weak *a priority* is implicit in the assimilation of logic to empirical theory.

This famous passage has been hailed as the classic formulation of an *anti-exceptionalist* conception of logic: in so far as its epistemology is

[1] This characterisation is, we are well aware, rough-hewn and open to objection. Its refinement is not on our agenda here.

[2] Quine (1953, p. 43).

concerned, Quine is saying, logic is no different from empirical theory.[3] Through Quine's vast influence, we have become used to this type of view of logic, inured to its shock value; but it is shocking. How could it possibly be right to claim that such propositions as –

(Identity)	If P, then P
(Universal Instantiation)	If (∀x) Fx, then Fa
(Modus Ponens)	If P and (if P, then Q), then Q

or their corresponding inference patterns – might count as *refuted* — that is, shown to be false, or unsound — merely by the consideration that rejecting them would help with 'simplifying quantum mechanics'? Of course, Quine may have meant to be appealing here to his background epistemic pragmatism, according to which the only types of good reason for belief that there are, are pragmatic reasons, not *epistemic* ones: what matters in the way of belief is what it 'pays' to believe. If so, those of us who wish to keep the notion of a non-pragmatic epistemic reason will have been given no reason to embrace his anti-exceptionalism.

To be sure, we might find—indeed, occasionally have found— that some seemingly intuitively evident logical principles lead to paradox so that something needs to be revised. So, intuitive evidence about basic logic is not infallible. But paradox-driven revision of logic would be revision of logic on *a priori* grounds not on empirical ones.

How did Quine propose to make this counterintuitive view plausible? As readers of *Two Dogmas* know, Quine's first move is to try to undermine the idea of the analytic, and in particular the idea that you could get to be justified in believing something just by understanding it. Having done that to his satisfaction, he then considers how else logic might be justified. The alternative he proposes is that the justification of everything, logic included, is *au fond* empirical.

How might we make sense of that? – in just the way that Quine tries to do, by arguing that logic is just another – albeit central – aspect of the 'web of belief', and that everything in that web earns its positive epistemic status by being part of the best, simplest and most effective overall empirical theory of the world. Thus, just as 'All swans are white' keeps getting better confirmed as we observe more and more white swans and no non-white swans, so, too, does 'If all swans are white, then this swan

[3] The term seems to have been coined by Williamson (2007). Other prominent contemporary expressions of anti-exceptionalism include Maddy (2002), Priest (2014) and Williamson (2013). But it is our impression that the broad idea is still widely accepted among philosophers who have not published in support of it, indeed that it approaches orthodoxy in many philosophical circles.

is white'. And just as 'All swans are white' might need to be revised on the observation of some non-white swans, if pleading hallucination proves not to be a viable response, so, too, might 'If all swans are white, then this swan is white' if the overall theory runs into something recalcitrant to which the best response would be to reconsider unrestricted universal instantiation. In short, like all scientific theoretical generalisations and inferences, logical laws may be *confirmed* only by empirical experience; and they may be rationally *revised* on purely empirical grounds. We will call the first claim *the Empirical Confirmation thesis* and the second *the Empirical Revision thesis*.

No doubt mainly because of the well-known difficulties that confront the project of developing a satisfactory non-empirical epistemology of logic, Quine's idea has, as we said, been largely hegemonic. A notable example of its sway was Putnam's proposal, which will occupy us briefly below, that the empirical successes of quantum physics give us reason to reject the classical distributive principles for conjunction and disjunction.[4]

In 1974, Saul Kripke gave a seminar at Princeton in which he argued that Quine's anti-exceptionalist view of logic, including its incarnation in Putnam's proposal, depended upon a confusion of *logical belief* and *inferential practice*. If you paid attention to this distinction, Kripke seemed to want to say, you would find that the claim that *any* logical law or inference pattern, no matter how basic, could in principle be *rejected* on purely empirical grounds and an alternative be *adopted* in its place, is unsustainable.

We are fortunate now finally to have an edited transcript of a lecture Kripke gave at Pittsburgh in the same year, which lays out in some detail how this so-called 'Adoption Problem' afflicts the Quine/Putnam empirical conception of logic. The issue has been in the wind and in recent times there has been a significant upsurge of interest in it, led by the interpretation and discussion of it in Romina Padro's doctoral dissertation and subsequent papers[5] and in papers by Alan Berger[6], Suki Finn[7], Corine Besson[8] and Michael Devitt and Jillian Rose Roberts[9]. In what follows, we first offer a statement of the problem—we will call it the *Original Adoption*

[4] Putnam (1969)
[5] Padro (now Birman) (2015) and Birman (formerly Padro) (2023) in this volume.
[6] Berger (2011).
[7] Finn (2019).
[8] Besson (forthcoming)
[9] Devitt and Roberts (2023).

Problem (OAP)[10] —as it seems to have been understood by Kripke himself and by these commentators, and argue that, so understood, it depends upon natural but resistible assumptions. So much the better for anti-exceptionalism. But we will then use that discussion as a springboard for developing a different adoption problem, arguing that, for a significant class of basic logical principles, there is indeed a difficulty in seeing how they might be 'freely adopted', thereby vindicating something close to the spirit of Kripke's original claim. This first part of our argument will enforce a significant qualification of the Empirical Confirmation thesis. In the concluding sections of the paper, we will turn to the question, specifically, of the empirical *revisability* of logic, arguing that when proper attention is paid to the role of reasoning in theory revision, it does indeed emerge that anti-exceptionalism, in full generality, is untenable. The Empirical Revision thesis is false: some logical principles, or types of logical principle, must be accorded not merely a central place in the 'web' but an epistemically[11] exceptional status *vis-a-vis* natural science.

2. What is the adoption problem?

2.1 Preliminaries—rule-following and inference

We need to begin with the distinction between *believing* in the soundness of a logical principle and *using* that principle in your reasoning. The main examples of relevantly basic logical rules on which the above-mentioned secondary literature focuses are Universal Instantiation (*UI*) and Modus Ponens (*MPP*). We forfeit nothing important for present purposes if we represent them by these simplified cartoon versions:

> (*UI*) For any predicates, *F* and *G*, given a premise to the effect that: *All Fs are G*, you may infer: *This, if F, is G*.
> (*MPP*) For any propositions, *P* and *Q*, given premises to the effect that: *P*, and *If P, Q*, you may infer *Q*.[12]

What it is to believe such principles to be sound is, presumably, clear enough – or at least as clear as belief ever is. But what is it to have a *practice* of *inferring* with such rules?

[10] 'Original' to emphasise the contrast with the different problem we will go on to develop below. We should stress that the reasoning that generates the OAP which we will outline is a reconstruction from the secondary literature.

[11] Later we'll consider further whether the exceptionality we identify is indeed properly viewed as *epistemic*.

[12] In order to accommodate Harman's (1986) observations, we'd have to complicate these cartoon versions.

A common thought is that to have a practice of following a rule – any rule, and so in particular such a rule of inference – is not a matter merely of being disposed to conform to it, but requires some kind of cognitively real downwards connection from the content of the rule to that disposition. We can encapsulate this thought by:

> (*Guidance Rule-Following*) To have a practice of following a rule, R, is to be disposed to conform to R *as a result of guidance* by R.

Familiarly, however, many philosophers find this to be too rich a notion of rule-following and are content to construe having a practice of following a rule R as consisting merely in a certain kind —what kind, exactly, is for them to explain—of conformity to R:

> (*Conformity Rule-Following*) To have a practice of following a rule, R, is to be 'appropriately' disposed to conform to R.

We will not here need to take a stand on this issue.

Let us turn now to asking what it is to *adopt* a logical principle. Kripke and Birman make it clear that the notion of 'adoption' that is the target of their adoption problem is that of *free, explicit* adoption: 'free' in the sense that the adoptive agent is conceived as being under no kind of *antecedent* epistemic obligation to accept the logical principle in question merely in virtue of their understanding of its meaning; 'explicit' in the sense that adopting the principle involves consciously accepting a formulation of the principle and modifying one's inferential behaviour accordingly, as opposed merely to slipping without consideration into a practice of inference in accordance with it.[13,14]

The reader should notice, then, that as matters are being set up, a proponent of the free adoptability of logic is committed to Guidance Rule-Following, at least as far as the implementation of freely adopted inferential principles is concerned. The idea of free adoption of a logical principle, as we glossed it above, is hardly intelligible except in terms which envisage a thinker or thinkers as considering the principle, moving to accept it and then acting on its instruction.

One further preliminary. We will assume, as is plausible, that we may, without any serious misrepresentation, think of the content of rules in general as schematisable like this:

[13] We will be working within a broadly internalist epistemological framework, according to which you have an 'antecedent epistemic obligation' to accept a claim only if you have (possibly idealized) access to the source of that obligation.

[14] In the context of a critique of Quine, it's obvious why adoption has to be free in the stipulated sense; it's less obvious why it has to be explicit. We will return to these topics.

(*Rules-in-General*) If conditions of type C obtain, one may/must/mustn't φ.

Our special interest, naturally, is in the application of Rules-in-General to rules of deductive inference.

2.2 The original adoption problem

With these preliminaries in place, let's turn now directly to the question of the free adoptability of logical principles. Based on the sources cited above, we offer the following rough statement of what seems to have been the generally understood gist of the problem:

> *The Original Adoption Problem (OAP):* You cannot simply explicitly adopt a basic logical rule, and thereby for the first time introduce inferential practices in accordance with it, because in order competently to follow it, you will need *already* to be an (implicit) practitioner of it.[15,16]

It's important to emphasize once again that it is free explicit adoption that's at issue here, as explicated above. Although we will revisit this issue, it's hard *prima facie* to see how, in contrast, there could be a philosophical problem with the idea that a community might just slip, unguided, into an inferential practice— might just somehow be caused, without any consideration of an explicit formulation of a logical principle, to be 'appropriately' disposed to comply with it in their reasoning.

Likewise, any problem must depend on the idea that the explicitly adopted rule is to *contentually guide* inferential practice in accord with it — there will not be any relevant problem with the idea that an explicit adoption might just causally *trigger* an appropriate practice (miraculous though that might be). No relevant difficulty is in prospect for the idea that one might confront an explicit statement of a rule and thereby be *sub-rationally caused* to launch into a practice involving conformity with it that one did not have before.

[15] We note in passing that there is a reason for talking of 'adoption' rather than of belief or propositional acceptance. 'Adoption' allows for the case where the rule in question is conceived of as an imperative, rather than as a normative proposition. Although *Rules-in-General* construes rules as normative propositions, we do not mean to foreclose on an alternative imperatival construal. 'Adoption' is suitably non-committal on this issue.

[16] Cf. Kripke (2023, p. XX): 'This is the problem. If he did not already reason in accordance with the pattern of inference we call "Universal Instantiation", telling him that it was true would do him no good; he couldn't "adopt" it as a hypothesis, he couldn't use it as an auxiliary to derive further statements. If he was not able to make the simple inference, "All ravens are black, therefore, this raven is black", then giving him some super premise like "Every universal statement implies each instance" as another premise won't help him either'.

2.3 The argument for the OAP

Why might there be a difficulty in supposing that any logical principle could in principle be adopted in the stipulated sense? A natural thought is that Guidance Rule-Following will amount to *inferential* rule-following: that to follow a rule in a sense involving guidance by it is (to be disposed) to infer *from* a representation of the rule *to* a mandate — requirement, permission or prohibition—for a certain type of (mental) action in a particular situation or series thereof. Call this the *Inferential Model of Rule-Following*.

Here is a simple example. Consider the rule for castling in Chess, cast in the kind of template determined by *Rules-in-General*:

If: your King and one of your (initial) Rooks meet each of these conditions:[17]

(i) They have not moved from their starting squares in the course of the game so far;
(ii) All squares between them are empty;
(iii) The King is not in check;
(iv) The squares that the King passes over are not under attack, nor is the square where it will land.

Then: you may first move the King two squares towards the Rook, and then place the Rook on the square immediately behind it.

On the Inferential Model, a competent player who is following the Castling rule in a particular game context checks for the satisfaction of each of the four conditions (i) -(iv) and, should they be satisfied, *infers* across the conditional that it is allowable to Castle (which she may or may not then elect to do).

Pari passu: on the Inferential Model, for a rule like *UI* contentually to guide a thinker in her inferential practice will thus require that she perform (possibly tacit) inferences from a formulation of *UI*: she will need to instantiate the rule to specific situations in which she finds herself and then infer across its ingredient conditional to a mandate for what to do. To work with our cartoon version of *UI*: the thinker has the rule:

(*UI*) For any predicates, *F* and *G*, given a premise to the effect that: *All Fs are G*, you may infer: *This, if F, is G*.

[17] 'Initial' is meant to preclude the Rook resulting from a Pawn promotion.

She reasons: 'I have the premise that [All Swans are white]. *UI* tells me that given a premise of that form, I may infer any instance of the form specified — in this case: that would be [This, if a swan, is white]. So, I may indeed so infer that [This, if a swan, is white]'.

Well, as the alert reader will instantly have spotted, once under the aegis of the Inferential Model, we do seem rapidly to embroil both *MPP* and *UI* in an adoption problem since any adequate statement of either rule will be both (schematically) general and conditional, along the lines

> If you have premises of such-and-such a character, you may infer a conclusion of thus-and-so a character.

An application of either rule to generate guidance about a particular case will thus, on the Inferential Model, itself require both instantiation and detachment. So, on the assumption of the Inferential Model, it appears that the OAP, as stated above, should be acknowledged, at least for both *MPP* and *UI*.

2.4 What's the significance of the OAP for the Quinean conception of logic?

Let's grant, for the moment, that the argument for the OAP is sound. How does it make trouble for the Quinean conception of logic?

It's not immediately clear how it would bear on the Confirmation thesis. That thesis says, we may recall, that our belief that *UI* is sound, and our practice of inferring according to *UI*, accrue such justification as they may have from their performing well in on-going empirical theorizing, and by their not being appropriately held responsible for any snags that such theorizing may have encountered.

However, this thesis does not look as though it needs to be embroiled in claiming that we came to be practitioners of *UI* by *adopting UI* in the specified technical sense. For all that the Confirmation thesis aspect of the Quinean conception cares, especially as applied to our practice of being *UI-ers*, it need take no stand at all about how we *came to have* that practice: it may have been hard-wired into us, or we may have picked it up by immersion, or what have you. However we came to the practice, the Quinean can still assert that it earns its keep by not falling under the shadow of doubt in empirical theorizing.

By contrast, the Revision thesis looks to be more directly affected. Because a proposal to rationally *revise* some aspect of our existing logical practice requires explicitly identifying the culprit in question, and, typically, *suggesting an explicitly formulated alternative in its place*, whose adoption is then supposed to contentually guide a change in our

inferential practice, the alleged fact that some basic inferential practices could not be adopted would put paid to the Revision thesis, at least in full generality, and so to the Quinean conception of logic of which it forms a crucial part.

2.5 How might the OAP be resisted?

However, fortunately for the anti-exceptionalist, it looks as though the argument for the OAP can be resisted, at several places.

A relatively minor point, first. Even if the argument for the OAP, as stated, is assumed to succeed, it still leaves us some way short of uncovering a *general* adoption problem. On that assumption, it is true that both *UI* and *MPP* will be implicated in following any general conditional rule, themselves included. And maybe we need not reckon with any other kind of rule; maybe *Rules-in-General*, given something close to the above formulation, indeed captures the structural essence of any rule. But in that case, for all we have so far seen, *other* basic logical rules — those for negation, for example — may yet be freely adoptable in a context in which *UI* and *MPP* are already practised. (To be sure, even this limited version of an adoption problem is enough to make trouble for the Quinean Revision thesis, understood as holding in full generality.)

The other two problems for the OAP are not as benign. The first stems from the essential occurrence of the term 'already' in the rough statement of the problem given above. All that has been pointed out so far is that all following of general, conditional rules will, according to the Inferential Model, involve inference in accord with *UI* and *MPP* (or if those rules are purely syntactically individuated, with principles licensing instantiation and detachment in some form). Hence, following those two rules in particular, since they are each general and conditional, will likewise effectively involve instances of the very modes of reasoning that they sanction. But that is less than saying that inferential practices exemplifying those modes must *antedate* the explicit adoption of the rules. The needed bridging assumption is

The Already Assumption:

> A rule, given as an explicit statement, cannot contentually guide subsequent inferential practice unless the inferential dispositions which need to be exercised in applying it are *already* in place independently and in advance of receipt of the explicit statement.

However, why must the relevant inferential practices *already* be in place? Why should not receipt of an explicit statement of, for

example, *UI* go on contentually to inform *both* regular (first-order) applications of it *and* the higher-order applications putatively involved in one's recognition of the licence the explicit statement provides for any such regular first-order application? The existing literature does not, so far as we have been able to see, notice, still less address, this question.

A second vulnerability in the argument for the OAP is its reliance on the Inferential Model of rule guidance – the idea, that is, that when rules guide practice, they invariably do so by the agent *inferring* from a formulation of the rule what it calls for, and acting accordingly. Some form of guidance by the rule is stipulatively built into the notion of adoption; but why does the guidance invariably have to be *inferential*? Can we not make good sense of the idea that one may have one's behaviour *rationally regulated* by a general rule without this involving *inference* from a representation of the rule?[18] Call a notion of rule-following that conforms to this idea an *Immediate Contentual Regulation* model of rule-following.

This idea needs much more discussion, which would take us too far afield to attempt here. But two thoughts suggest that a proponent of the free adoptability of logical rules should look on it favourably.

First, even if we set aside the notion of *adopting* a rule of logic, given the assumption that deductive inference is itself an example of Guidance Rule-Following, to suppose that the guidance exerted by any rule of inference is itself essentially received by inference is to suppose that in order rationally to perform any particular inference, one has to perform another — one that consists in inferring that one has a mandate for the former. That makes inferring into a 'supertask'. We finite mortals cannot perform such tasks. That is to say, we know on purely general grounds, and quite apart from the present concern with the notion of adopting a rule of logic, that it is viciously regressive to combine both an Inferential Model of Rule-Following and a Rule-Following Model of Inference. One or other of these must go.[19]

Now, since the very idea of *adoption* (in the specified technical sense) of logical *rules* is committed from the get-go to a Rule-Following Model of Inference, any theorist who buys into that idea will have an iron-clad commitment to jettisoning the Inferential Model of Rule-Following, at least in full generality. So, any proponent of adoption will have to take very seriously the idea of non-inferential guidance by a rule.

[18] As we read her, a version of this suggestion is canvassed in Besson (forthcoming).
[19] For further discussion see Boghossian (2003) and Boghossian (2008).

Second, we seem to need some such non-inferential notion of acting upon a policy or general intention in a wide class of quotidian contexts. Consider a tennis player who, in connection with an upcoming contest, forms the conditional intention to advance to the net on his first serve but stay back on his second. It seems as though he should be capable of repeatedly acting *on* that intention without constantly having to infer what it calls for in the course of the hurly-burly of the match. Or consider someone who, frustrated by the difficulty of staying on top of email that has drifted down in her inbox, forms the conditional intention to answer immediately upon receipt any email that calls for an answer. It would seem as though adoption of that Email Rule could regulate her behaviour, leading to her indeed answering any answer-demanding email immediately upon receipt, without her constantly having to infer that that is what is called for by her general policy.[20] How does the OAP dialectic shift if we move to an Immediate Contentual Regulation model of (some) rule-following and of (some) inference in particular?

Well, the manner in which an explicitly adopted rule is to guide inferential practice need not now involve inference at all. It can involve instead (mental) action that is contentually regulated *by* the rule but does not involve inference *from* the rule. So, of course, there cannot be any implication that it must involve inference in accord with the very rule(s) concerned.

If, therefore, we buy into the suggested idea of non-inferential contentual regulation, we have the resources to block the reasoning to the OAP without surrendering the idea that genuine rule-following, including the following of explicitly adopted rules of inference, involves guidance by rule.[21]

Does that dispel any objection to the free adoptability of any logical law or inferential pattern?

2.6 A new adoption problem

It does not. An issue remains. Consider Padro's example of the fictional Harry.[22] Harry does not yet have the practice of universal

[20] It's beyond the scope of this paper to go into detail into how to make sense of non-inferential guidance by a general rule; for an attempted account see Boghossian (ms).

[21] Although it may impress as mysterious, we *must* buy into that idea if we want to allow the coherence of the notion of an ideally rational subject whose every inference is a fully self-conscious mental act. Of course, some may regard it as debatable whether we should want to allow that, or whether instead what should be admitted is that any rational subject who infers at all must perform some inferences, properly so called, to which the idea of guidance by rule is inapplicable.

[22] Birman (2023, p. XX).

instantiation. He is not disposed to infer 'This F is G' from an acceptance of the statement 'All Fs are Gs.' He is supposed to have *some kind* of understanding of the universal quantifier, but not one that commits him to the rationality of being a *UI* inferrer.[23] The question is whether he can rationally come to appropriately extend his practice to becoming a *UI*-inferrer by receiving a suitable statement of the rule of *UI,* along with some empirical (Quinean) reasons for believing the rule to be useful.

Can someone whose understanding of the quantifier so far gives him nothing in virtue of which he might recognize a commitment to the practice of *UI*, freely adopt *UI*, that is, explicitly take it on board as a new rule, and thereby institute a practice of inference in accordance with that pattern?[24]

Well, how is Harry to understand any particular general statement of *UI* that he may be given? Such a statement will itself embed (something equivalent to) universal quantification. So, if Harry parses it in terms of the *instantiation-neutral* understanding of the quantifier that he has before he gets the statement, then what he thereby understands is nothing he can use to derive and practise the new pattern of inference that the statement is intended to mandate.

Thus, if an understanding of the explicit statement of *UI were* to somehow prompt him into the appropriate extension of his inferential practice, this 'prompting' cannot be viewed as contentual guidance — as teaching given to him just by the statement of *UI* itself —but must have worked in some other way so as to trigger a shift in his understanding of 'all'.

So, we seem to have unearthed, even after we discard the Inferential Model of rule-guidance, a different but potentially interesting form of

[23] An anonymous referee for this journal questioned whether these two stipulations about Harry are not in tension with one another: could someone have *any* grasp of the universal quantifier without some tendency to perform universal instantiations? We do not think the stipulations are in tension with each other. Any suggestion to the contrary would amount to the claim that only someone who grasps the elimination rule for it can be credited with *any* kind of understanding of the universal quantifier. We agree that Harry is in an extremely bizarre conceptual state, but it seems at least behaviourally intelligible that someone could master a range of cases of the *introduction rule* for 'all', without having the grasp that Harry is said to lack. For example, Harry might do reasonably well with judgements based on observation like, 'All the biscuits on the tray are chocolate fingers'. Casalegno (2004) explicitly proposes such a possibility, and Williamson (2020) defends Casalegno on this point.

[24] If *UI* is already implicit in his prior understanding of the quantifier, his explicit acceptance of the rule will not be a free adoption in the intended sense, but rather the coming to a piece of knowledge of a prior commitment.

adoption problem. The argument for the OAP was that, for *UI* and *MPP* at least, free adoption is ruled out because adoption of these rules—taking them explicitly on board and applying them in inference —would require, so the argument ran, that one already practise in accordance with them in inference. Our new point, by contrast, is nothing about what inferential dispositions need to be in place 'already'. It is that, however 'contentual guidance' is supposed to work, the understanding that Harry has of the principal logical operator in a statement of a given inferential rule can only guide him towards a practice in conformity with *it*—that is, with his prior understanding. Accordingly, he cannot be contentually guided towards any *change* in his practice by the explicit statement.

2.7 Generalising and qualifying the new problem: the lambda operator adoption problem (LOAP)

To fix ideas, consider a case where a community of agents who have some notion of a logical operation, λ, fail at present to infer in accordance with what we regard as a basic rule of inference for it, R. Are they in position *freely* to *adopt* inferring according to R?

Once again, it's vital to bear in mind the relevant idea of *free* adoption. As the Quinean/anti-exceptionalist intends, or ought to intend the notion, freely adopting an inferential rule R involves coming to perform R-inferences, not because one recognizes that one is antecedently committed to, or justified in, falling in with this pattern in virtue of one's prior understanding of the form of inference and logical operators deployed in R, but because one comes to regard so inferring as an independently good idea (motivated, perhaps, on the basis of some empirical consideration).

In order to be available for free adoption in this sense, therefore, R cannot be *epistemically analytic* of λ, as understood by the adoptive agent, in either of the two ways recently distinguished by Boghossian.[25] It cannot be that the agent's antecedent understanding of λ, together with grasp of the syntactic form and of the remaining constituents in the statement of R, *constitutively requires* a disposition to infer according to R. Nor can it be that the agent's antecedent understanding of λ, together with grasp of the syntactic form and the remaining constituents of the statement of R, although not constitutively requiring a disposition to

[25] See, for example, Boghossian (2017).

infer according to R, provides a sufficient rational *basis* for the agent's acceptance of a practice of inferring according to R.[26]

Both of these cases contrast with *free adoption* since, in both cases, the pattern of inference in question is antecedently mandated by one's understanding of λ. 'Free adoption', remember, is to be something that people may do whose antecedent understanding of λ *underdetermines* acceptance of R. However, when — as was the case with Harry and *UI* — the very statement of R makes essential use (not mention) of the operation λ itself, the adoptive agent's understanding of the patterns of inference licensed by R will be constrained and limited by their antecedent understanding of λ.

The resulting dilemma is simple. In order to be available for free adoption, R must be a *new* rule, independent of one's antecedent grasp of its principal operator, λ. But if all one gets by way of explanation of the new inferential practices it 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. Thus, there is indeed an adoption problem for a wide class of logical notions and associated rules. And, in contrast with the OAP, it is a problem arising independently of the Already Assumption and the Inferential Model of rule-guidance.

How wide a class? Let us say that a logical notion is *basic* just if a canonical statement of its characteristic inferential profile must essentially involve (metatheoretic) use of (a counterpart of) λ itself. We shall henceforward term the expression for any such basic logical notion a *lambda-operator.* The preceding applies to any lambda operator.[27]

Thus, the point, at least potentially, afflicts all the standard sentential logical connectives and the existential and universal quantifiers.[28] For any such basic logical notion, λ, any attempted such canonical statement which is to non-conservatively extend the relevant agents' 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

[26] Arguably, for instance, grasp of the comparative 'taller than' doesn't constitutively involve recognizing that the relation is transitive. But it is plausible that a person who understands 'taller than' can, just by drawing upon their prior understanding of it, come rationally to recognise its transitivity and so infer accordingly.

[27] We have obviously appropriated the term 'lambda-operator' for our own purposes; it should not be confused with the notion of a 'lambda function' familiar from the theory of computation.

[28] There is a certain relativity here, depending on whether a certain operation is or is not treated as definable in terms of others. Disjunction, for example, is basic in intuitionist logic but, if negation and the conditional, for example, are treated as primitive, not so in classical logic.

which *already* diverges from that possessed by its recipients and is thus not available to them at the crucial point. Or if it is available to them, then that understanding has already to be implicit in their understanding of λ as it existed before the adoption, so that what we are concerned with is then, epistemically considered, a reflective *extrication* rather than an adoption.

It might be thought that this argument overreaches: that, in the specified sense, there are no lambda operators. For is it true that any explication of something like the rule *UI* will essentially involve 'all'? Couldn't the basic point behind a rule like *UI* be captured, rather, by

$\neg\exists x\neg Fx/Fa$?

Certainly. But now we would be taking as basic negation and the existential quantifier, and so those would count as lambda operators.

We can safely assume that within any given language or idiolect, some connectives will be semantical primitives while others may be definable in terms of them. If all of ¬, ∧, →, ∀, and ∃ are semantical primitives, then they will all count as lambda operators in our sense. If only some are, then only those will count. Thus, our notion of a lambda operator is relative to the choice of semantical primitives, a choice that, in any language, will have to be made.

In summary: if λ is basic, any adequate explication, S, of its inferential profile, whether it, for example, proceeds by a characterisation of the truth conditions of statements in which λ is the principal operator, or by the articulation of (introduction and elimination) rules of inference for λ, will have to make use of the concept of λ itself. In order to adopt such a statement, S, as regulative of her practice, an adoptive agent must first understand it and must therefore have *some* prior understanding of λ. If that understanding already embraces the inferential practices underwritten by S, S is not available for *free* adoption. If it doesn't, then what is supposed to be novel about S and the practices it sanctions cannot be expected to be appreciable by the adoptive agent.[29]

[29] An objection needs treatment here: many, including Wright and Hale (2000), have supposed that the explicit stipulation of the truth of a sentence or sentences containing a previously undefined or only partially understood notion, can somehow amount to/achieve a fuller explanation of it. If this can indeed occur, is acknowledgement of that at odds with the thrust of the dilemma just argued? We don't think so, because such implicit definition, when it can be successful, draws only on the understanding (i) of *other* notions besides the definiendum featuring in the stipulated sentences, and (ii) of their syntax, whereas in the cases we are concerned with, for the stipulation to have the intended effect, some sort of prior understanding of the definiendum itself, ⌊, is required.

3. The Empirical Confirmation thesis

3.1 A moral of the LOAP: the acceptability of basic logical principles cannot have an exclusively Quinean justification.

As we noted at the start of this essay, the Quinean picture of the epistemology of logic entrains both an Empirical Confirmation thesis and an Empirical Revision thesis. Let us look now at the first in the light of the LOAP as now developed.

Say that a justification for an accepted hypothesis, H, of an empirical theory, is *Quinean* just in case it consists in the discovery that the theory performs well (by whatever criteria for the theory performing well that the theorist deems appropriate) before the 'tribunal of experience', and would work less well if deprived of H. Quine's idea, and that of the contemporary anti-exceptionalists, is that the logic, including the basic logical principles, which an empirical theory incorporates, insofar as it admits of justification at all, admits *at most* of Quinean justification in this sense.

One moral of the LOAP, we suggest, is that this claim is unsustainable.

For concreteness, let us continue to work with the example of *UI* and the predicament of Harry, although everything we say here applies, for any particular logical language, to all lambda operators and all basic rules for them. Harry, recall, has some kind of understanding of 'all' but, so far, no practice of universal instantiation. There are two cases to consider.

First suppose that *UI* is not a candidate for *free* adoption by Harry, but rather that inferential practice in accord with *UI* is indeed built in to his antecedent understanding of 'all', rather as the admissibility of inference via transitivity is built into an ordinary understanding of 'is downstream of'. However, it simply hasn't occurred to him that this is so. Then intuitively Harry already has a justification for adopting *UI* — the same species of justification that someone has on whom it dawns, thinking about it for the first time, that 'is downstream of' is a transitive relation and accordingly adopts a principle to that effect. So the anti-exceptionalist claim, that the *only* justification one can have for a basic logical principle is Quinean, cannot be sustained in this first case.

Suppose instead that *UI is* a candidate for free adoption by Harry: nothing in his prior understanding of 'all' mandates his accepting it. Can he acquire a Quinean reason to do so? It seems not. It is no good his just spontaneously accepting what we will regard as an adequate formulation of *UI* and seeing what pay-offs accrue from working with it. There won't *be* any such pay-offs since he won't be disposed to apply the

principle in the way that will be needed in order to generate them, but will understand it in the same old 'instantiation-neutral' way as before. For the same reason, it is no good our simply telling him,

> (*Benefit*+) If you fall in with this extension — (and here we offer him a characterisation of proper *UI*-practice)—of your erstwhile pattern of use of 'all', you may expect the following benefits.....

since, assuming 'all' to be a lambda operator, the characterisation we thereby offer him of the recommended extended practice will, if it is to convey what we want, presuppose a kind of understanding of 'all' which, by hypothesis, he doesn't have. The same limitation on uptake of a characterisation of a recommended extended practice will go, *mutatis mutandis*, for any lambda operator.[30]

3.2 Epistemic analyticity and Quinean pragmatism

If this argument is cogent, a strong anti-Quinean conclusion might seem to be in the offing: certain basic logical principles must indeed be acknowledged to have an exceptional epistemic status when compared with the general run of empirical hypotheses — namely, you cannot have *only* Quinean (empirical) reasons for adopting them. Indeed, if you have reason at all, it must include the kind of reason we suggested is operative with the transitivity examples like 'is taller than' and 'is downstream of': the kind of reason, that is, implied by the idea of epistemic analyticity.

But now, supposing that is accepted, how damaging, really, is this conclusion to the Quinean/anti-exceptionalist? It might be thought 'not very', for two distinct, though related, reasons.

First, the anti-exceptionalist may reply, to allow that certain basic logical principles are epistemically analytic relative to a given concept of the operator they feature does not provide those principles with a real *justification*, but only with something more like a rationale: we can see why it might make sense for someone using the concept in question to employ the inferential rule it licenses, but this is consistent with the concept in question being no better than — for the sake of an egregious example —'tonk'.

Second, even if we allow that the lambda operator in question is, unlike 'tonk', a legitimate inferential concept, it may seem that there might still be a question whether the concept in question is an *optimal* one, or even a good one, for a thinker to use in empirical theory.

[30] Thus, in contrast with the OAP, the LOAP does have a bearing on the Empirical Confirmation thesis of the Quinean picture.

The first question seems hardly worth taking seriously, for it is arguable that 'tonk' expresses no real concept, since its constitutive rules are not coherently followable. By contrast, 'if' and 'all' not only seem to be perfectly fine concepts; they, or something very like them, are arguably *indispensable* for any thought more sophisticated than the most elementary. In any case, it borders on the incoherent to suggest that, strictly, we can't know that 'If P, then P' holds true on the basis of understanding 'if', but can only know the qualified version that 'If P, then P' holds true *if* 'if' expresses a legitimate concept in good standing.

The second challenge, however, might seem more credible. Perhaps it is out of the question to doubt whether 'if' or 'all' are legitimate concepts. Still, isn't it consistent with that to wonder whether they are the *optimal* concepts with which to think?

To illustrate this point, suppose, as is in any case plausible, that the basic geometrical axioms of Euclid and Riemann, configuring notions like 'point' and 'line', are epistemically analytic of their *respective* versions of those notions, so that Euclidean and non-Euclidean geometries do not, strictly speaking, disagree on any geometric claim. It would not follow from this supposition that our views on the appropriate geometry for reasoning about physical space are not subject to empirical confirmation or revision. All that would follow is that the question which geometry correctly describes space would be equivalent to the question which set of basic geometrical concepts should best be used to characterize physical space.

Similarly, it might seem, even if some basic logical rules are epistemically analytic of the operations they configure, it would not follow that we could not raise the question: which logical concepts are most advantageous or otherwise appropriate in reasoning about the world? — and then move to abjure certain logical principles that were epistemically analytic of, for example, suboptimal notions of the conditional or the universal quantifier.

But there are several difficulties with this line of thought. In the first place, it is not clear that the requisite analogy between logic and geometry can be sustained. In the case of geometry, we can make sense of the difference between pure and applied geometry, distinguishing between geometric primitives that define abstract geometric structures and those that correctly characterize the structure of physical space. What, though, would constitute an analogous distinction in the case of logic? What's the difference between pure and applied logic? If you are thinking with classical constants, double negation elimination is valid

both for thought and for 'the real world'; and if you're thinking with intuitionistic constants, it is just as invalid for both.

In any case —secondly—what does 'works optimally well in scientific theorizing' have to do with being a *valid* logical principle? The validity of a logical principle licenses its use in arbitrary reasoning about counterfactual circumstances, including circumstances in which physical law is hypothesised to fail. How exactly does the entrenchment of principles of inference in empirical theory that has proven to work well for our *actual* circumstances rationally mandate confidence in their use in reasoning about circumstances of quite another kind?

That concludes what we want to say here about free adoptability and the Empirical Confirmation thesis.

4. The Empirical Revision thesis

4.1. Types of revision

As we noted at the beginning of this paper, it is consistent with supposing that basic logical principles/inferences have justifications that are *a priori*, that they can be *rejected* on purely empirical grounds. Weak *a priori* justification does not entail strong *a priori* justification. We turn now, then, to a discussion of the Revision thesis. Does our Lambda Operator Adoption Problem have any bearing on this second, all-important component of Quine's conception of the epistemology of logic?

It bears very directly, of course, for exactly the reasons already discussed, on any scenario in which a putatively empirically motivated rejection of an accepted logical law configuring a lambda operator is combined with a proposal that we *adopt* a different law for the operator in question in its place. But there are two other scenarios to consider: that in which (call this *Mere Rejection*) we have a proposal simply to *reject* on empirical grounds an accepted law, *without replacement;* and that in which (call this *Restriction*) what is proposed is rather a limitation of the scope of the original law.

Historically, the most notorious example of the Quinean conception of logic in action seems to have been an example of the second scenario: Putnam's (1969) proposal that some of the singularities of quantum mechanics would disappear if we rejected the classical conjunction-disjunction distributivity principles:

$P \wedge (Q \vee R) \leftrightarrow (P \wedge Q) \vee (P \wedge R)$
$P \vee (Q \wedge R) \leftrightarrow (P \vee Q) \wedge (P \vee R)$

in cases where 'P', 'Q' and 'R' are propositions whose truth-values are subject to quantum indeterminacy. In putting his proposal forward, Putnam never, so far as we are aware, advocated dropping the distributivity laws across the board. His idea was rather that we should think of quantum phenomena as representing an exception to their scope.[31] More generally, his thought seems to have been, we should be open to the idea that logical laws may be subject to contingent limitations on their scope, contingent limitations which empirical enquiry may disclose.

4.2 The LOAP and the Empirical Revision thesis

We will argue that neither Mere Rejection nor Restriction can be true in full generality. Some of our arguments will be based on considerations adduced in the course of developing the LOAP; others will be based on related, though distinct, considerations deriving from the role of reasoning in any attempt to make sense of the rational revision of basic logical principles.

Let's begin with Mere Rejection. It is initially hard to see how an *adoption* problem can automatically make trouble for any proposed revision of logic that consists merely in a suggestion to *reject* some basic rule. Any form of putative adoption problem, cogent or not, raises a difficulty for the idea that we might extend our inferential practice with a given operator by extending the repertoire of inferential principles governing that operator that we accept. No such process is involved in recommending rejection of a particular pattern of inferential behaviour.

However, that is not to say that Mere Rejection is off the hook. To appreciate why, suppose this time that Harry, now an adult, accepts both *MPP* and *UI* and practises them perfectly normally in the context of an

[31] Actually, it is moot what exactly Putnam is best taken as having been recommending in his famous paper. One take on his suggestion — at odds admittedly with his own usual gloss on it —would render it not as that, when reasoning about quantum-phenomena, we should continue to work with notions of conjunction and disjunction, properly so termed, but free them from the mutual ties codified in the classical distributivity laws, but rather that we should set those notions aside in favour of a pair of *ersatz* operations analogous to them in some respects but stopping short of the validation of those laws (specifically, that we should think in terms of a pair of new operations whose interpretation would be given by the operations of meet and join on a non-distributive lattice). This interpretation would, superficially viewed, assimilate Putnam's proposal to a further example of the kind of situation displayed by the Euclidean and Riemannian conceptions of point, straight line and so on, as mooted above. The huge problem, of course, is whether, in the case of either quantum-theoretic *ersatz*, we have indeed any concept of such an operation, that is, whether there *is* any intuitively thinkable binary connective taking one from a pair of constituent thoughts to a determinately intelligible compound thought in which they are adjoined, which comports itself in the manner Putnam requires.

We are grateful here to Tim Maudlin.

empirical theory which is repeatedly hitting snags. Might he rationally treat these snags as an empirical reason for ceasing so to practise? If so, the reason must, presumably, consist in his considering that he has accumulated forceful evidence for a proposition like

> (*Benefit-)* If I fall in with the total rejection of *UI*- or *MPP*-practice, I can expect to avoid the problems and singularities I have been having.

Well, it is obvious enough that the implementation of *Benefit*- in Harry's future practice — his acting on the recommendation it offers —will require him to instantiate it to particular cases and to detach across its contained conditional. So if total rejection of *UI/MPP* is what is recommended by *Benefit*-, Harry won't be able coherently to rationalise actions upon his new policy.

To be sure, if the recommendation of a *Benefit*- type proposition were merely that we should *partially* restrict *UI/MPP* practice, the instantiations/detachments which actioning the recommendation required might be among the permitted exceptions — we'd need to see a worked example to be sure how that might play out. But the foregoing is still enough to show that Quinean Mere Rejection cannot be true in full generality: the proposal that we might have rational reason to reject *UI* or *MPP in toto* is off the table.

The situation with Restriction is more subtle. While we foresee no absolutely general difficulty with the idea that empirical reason might emerge to recommend that we trim an inferential practice in certain respects, partial restriction does raise some special issues in our two central cases: the conditional and universal quantification.

First, consider the latter. Adding restrictions to a rule still leaves its generality uncompromised at a metatheoretic level. Restricting the rule involves introducing more conditions into its antecedent. Suppose we start with:

> Given (any instance of) $\{A_1 \ldots A_n\}$, you may deduce the corresponding instance of P,

but then we are tempted to restrict this to:

> Given (any instance of) $\{A_1 \ldots A_n\}$, *and provided these premises meet such and such additional (maybe syntactic) conditions,* you may deduce the corresponding instance of P.

Then the point is that *this* restricted rule is still unrestrictedly general — it says that whenever its antecedent conditions are met — restrictive

though they may be — the deduction of the corresponding P is mandated. Someone who applies the restricted rule is arguably still bringing to bear unrestricted instantiational and detachmental modes of thought at a metatheoretic level in so doing.

Call this point the principle of the *Generality of Restricted Generality* (GRC). Its immediate effect is that while we can impose syntactic or other restrictions on the use of *UI* in a specific object language, we cannot coherently impose such restrictions at every level of our thought.

There is a related point with the conditional. We might indeed impose a restriction on, say, the premises for *MPP* — for instance, that the transition from P and $P \rightarrow Q$ to Q is allowed only if P meets some additional condition, C.[32] But then we are allowing a new out-rule for the conditional:

$$\{P,\ C(P),\ P \rightarrow Q\} \Rightarrow Q,$$

to be unrestrictedly valid. There seems to be no way past the point that, whatever the detail of its formulation, our thought, if rational, must at some level comply with some unrestricted principle of detachment.[33]

4.3 The rational revision of logic and its role in reasoning

We now wish to canvass an alternative line of argument against the Quinean claim that no inferential principle is immune from rational revision on empirical grounds. The argument will turn on paying close attention to the role of logic in the meta-reasoning that will be needed to make a case for the rational revision of any logical principle. This line of argument, too, is anticipated in Kripke's material from 1974, although it should be sharply distinguished from the 'adoption problem'. Kripke formulates the point thus:

> The point is that 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

[32] For instance, one not uncommon reaction among students first encountering Vann McGee's purported counterexamples to Modus Ponens is to suggest that the principle guarantees sound inference only in cases when the minor premise is not itself a conditional.

[33] We suggest that these considerations block the 'Neurathian' thought, put to us by an anonymous referee for this journal, that while whenever we make revisions to our logic, we must hold some general principles fixed above the 'Quinean fray', no particular principles need to be held fixed once and for all, and that any particular principle may, after enough rounds of restriction and revision, be gradually whittled out of existence.

Not so, we say. Instantiation and detachment, in some respective forms, will and must survive any amount of 'whittling'.

> logic in order to draw these consequences. There could be no neutral ground in which to discuss the drawing of consequences independently of logic itself. This is the basic point that I want to make.[34]

What follows develops this basic insight. Suppose we have somehow arrived at a plausible (by whatever criteria) formulation of a logical theory, L, which we intend as a formalisation of the pre-theoretic patterns of inference that we subscribe to in our practice of a particular empirical theory, T. And suppose we run an experiment whose result is *prima facie* disconfirmation of some prediction of T. As good Quineans, we want to reserve the option of redirecting this disconfirmation against some element, R, of L — an axiom or a rule of inference —used in deriving the relevant prediction from T. How are we to make out that this is a rational option?

We need to represent the situation as one in which R features amongst the range of things that are collectively *responsible* for the rogue prediction. That requires that we make a judgement to the effect that the rogue prediction follows from T when R is accepted in the underlying logic (and perhaps is underivable otherwise). That is a *metalogical* judgement. Such judgements can be quite complex and verifiable only by explicit reasoning about (the form of) the premises provided and the capacities of the inferential machinery concerned. And this meta-reasoning will itself draw on a set of principles of inference, whether explicit or merely tacitly accepted. The subjection of T *along with* L to simultaneous empirical test will thus involve reasoning mediated by a background set of principles of metalogical inference, L*.

So: what is to be the system of metalogical inference, L*, in which the simultaneous testing of L with T is to be thus conducted? How does L* relate to L? To be sure, we just now introduced L* as a *meta*logic, a logic in which we reason about what follows from what in L. But that is a difference between the ranges of statements to which, in the scenario we are concerned with, L and L* are respectively being applied. Logical principles are not parochial in scope. It will still make sense to ask whether any of the principles that feature in L* — in what we consider to be valid meta-reasoning about the deductive potentialities of L — also feature in L itself.

Ideally, L* would be disjoint from L, and would be independently highly credible. Only then would it offer a medium for properly

[34] Kripke (2023, p. XX).

independent testing and assessment of L in the light of the failure of the predictions from T enabled by L. But this ideal looks unrealistic. If L is, for example, a propositional and quantificational logic of any normal degree of strength, L* will surely significantly overlap with it. *How can we test the empirical credentials of a logic by using what is essentially the same or a substantially coincident logic in reasoning about it*?

Here is where the distinction —emphasised by Kripke in the passage we quoted earlier —between reasoning and an explicit system of logical laws begins to matter. Suppose we deduce using L a contradiction between elements of T and observational data, O, and are tending to the view that it is not T, but L that is to blame— specifically, that the rogue may be one of L's basic principles, R. In order for this suspicion to be a rational option, *we have to remain confident that the original derivation of the contradiction is good*, that is, that we really have shown that the deductive resources of L do indeed generate a conflict between T and the observational data. So we had better have no qualms about the good standing of the L*-principles essentially involved in that demonstration. What if a counterpart of R *itself* is so involved?

Casually regarded, the situation might seem merely to be a form of *reductio* of R. But that reaction is not dialectically stable. Rather, if R is indeed essentially involved, not as a premise, but as a rule of inference in the reasoning whereby we obtain the contradiction, treating the situation as a *reductio* of R plunges us into a state of aporia. For in order rationally to justify a proposal to revise R on the basis of the contradiction, we need to repose trust in the consequence relation for L* — to trust that L*-consequences *really are* consequences; and in order to trust in that, we need to trust that conclusions licensed by R, as one of the L*-principles involved, *genuinely follow*. So we wind up both distrusting R — our proposed Quinean conclusion — and committed to trusting it — otherwise we cannot justify the distrust. R becomes, as it were, an unreliable witness to its own unreliability.

If this is right, then we should reject the Quinean thought that when one is led to a situation of 'recalcitrance' in the practice of a scientific theory, it can in principle be rational to hold to account not merely the empirical scientific premises, but *any* aspect of the inferential machinery involved. There is incoherence in the idea that the case for revising a rule of inference, R, might rest on a derivation of a problem— a situation of 'recalcitrance'—in circumstances where recognition that you

have a problem at all itself relies essentially on a metatheoretic use of R.[35] The key thought is that the belief that *you really have a problem*, which rationality requires you to remedy, must rely on the belief that your derivation of the contradiction, or whatever the problem is, is sound, so on a belief that R is good. If you then query that, you undermine your reason for thinking that you have a problem in the first place.

It may be questioned (and was by Joshua Schechter, in conversation) whether in order to draw the conclusion that R is unreliable we need, in the way suggested, to presuppose R's reliability. Is not the situation rather like that where a witness, Fred, (we can add: under oath) says, 'I am unreliable'? We exactly *don't* need to trust Fred in order to rationally accept what he has said. Rather we can reason by dilemma. Either he is reliable or not. If he is, we can accept what he has said, so he isn't; and if he isn't, well then he isn't. Either way, he isn't reliable. Likewise with R: if it is sound, then the derivation of the contradiction shows there is a problem with assuming it to be sound; and if it isn't sound, then it isn't sound. So either way, we can rationally take R to be unsound.

But there is no relevant parallel. Here's why. The situation with R will be complicated by the fact that in any actual case, there will be, for Quinean reasons, a *holism* of factors contributing to the apparent problem besides R. True, we presented the case as involving some strong background reason why R is *prima facie* supposedly especially suspect, but even so, the Fred analogy becomes questionable when we reckon with the holism involved in any R-case. For the holism entails that we *don't have* the conditional we need for the first horn of Schechter's dilemma, namely, 'If R is sound, then the derivation shows there is a problem with it'. Rather we should think instead that if R is sound, then the problem unearthed by the derivation is not to be attributed to R but must after all be attributed to some other factor(s) in the holism. On the first horn, we need to reassess our judgement that R is the prime suspect in the genesis of the problem. So the dilemma collapses.

4.4 Avoiding the R-pickle—the exceptionality of detachment and instantiation

The illustrated pickle — let's call it the 'R-pickle'—doesn't arise if the L*-principles involved in the original demonstration of the problem happen not to include any counterpart of R. It is of interest that there are

[35] We needn't worry what exactly the nature of the problem is— above, we worked with the toy idea of getting results inconsistent with observation, but that was just to fix ideas; the key point is only that the problem be arrived at by inference.

ways to ensure that they don't— to an extent. Many more senior readers will recall that before the more recent vogue for teaching elementary logic by reference to systems of natural deduction, older textbooks would outline axiomatic systems. The rules of inference involved in the latter were usually pretty sparse. Where for example a natural deductive exposition of classical sentential logic might include a primitive rule of disjunction elimination, an axiomatic presentation might include an axiom like

$$((P \wedge R \rightarrow T) \wedge (Q \wedge S \rightarrow T)) \rightarrow (((P \vee Q) \wedge (R \wedge S)) \rightarrow T),$$

whose application in proofs would then be left to rules of admissible substitution for the sentential letters involved, and, of course, to *MPP*.[36] So to avoid the R-pickle in any particular case, we can always schematically axiomatize the relevant R, add it as an additional assumption to T, and fall back on an underlying metalogic containing just substitution rules and *MPP*. So L* can indeed always be quite a bit sparer than L may be.

'Quite a bit', but only provided, once again, it contains resources for the OAP's original suspects, detachment —*MPP*— and instantiation — *UI*. Casting rules of inference in the form of conditional schematic axioms accomplishes nothing unless we have an out-rule in L* for the conditional, tantamount to *MPP*, and rules for the legitimate instantiation of schemata that are arguably epistemically tantamount to special cases of *UI*. To attempt to call *these* rules to account before the tribunal of experience must, it seems, be to succumb to the R-pickle.

As, in effect, was noted earlier, the hard-line anti-exceptionalist has a little wiggle room. They might seek to formulate a restriction such that, in the undoubtedly far-fetched kind of a scenario where *MPP* or *UI* are reasonably regarded as under empirical pressure, it is only *unrestricted* versions of them that are under suspicion of spawning a problem, and only r*estricted* versions that are needed in the L*-reasoning on which we need to rely in order to be rationally confident that there is indeed a problem to be dealt with. We are happy to leave to the anti-exceptionalists the project of persuading us of plausible examples of how this might work. When they are done, we will be able to point out that, for the reasons noted earlier, they have still retained unrestricted versions of instantiation and detachment.

What we may conclude is that *some versions of MPP* and *UI* respectively must indeed be treated as exceptional by any coherent

[36] Pedagogically, this kind of thing placed great (often unmeetable) demands on the substitutional insights of even very good students.

management of sufficiently rich systems of empirical belief. Principles of detachment and instantiation may be 'tweaked' but not abandoned wholesale. But the reason is not, as the Original Adoption Problem had it, that one must already be a practitioner of these principles in order to put them into effect. Rather the overall patterns of reasoning licensed by these principles cannot be regarded as challenged outright and in full generality by experience except at the cost of the cognitive incoherence of the R-pickle.[37]

This should certainly be regarded, by the anti-exceptionalist's lights, as a kind of exceptionality. Is it properly regarded as a kind of *epistemic* exceptionality? We suggest a qualified answer, 'Yes, but...'. 'Yes', because the argument, if sustained, has the effect that the special place occupied by principles of detachment and instantiation in our thinking cannot, as on the Quinean picture, coherently be credited to their long-term success in empirical theory and consequent entrenchment. The basis of our knowledge of their good standing is nothing like that of our knowledge of empirical theory based on generations of accumulated evidence. The qualification, 'Yes, but...' is merited because the argument from the R-pickle is silent about the epistemology of these fundamental modes of inference: if we do indeed know that they are sound, no positive account is suggested by the R-pickle argument alone about the provenance or character of that knowledge.

4.5 Coda: a further anti-Quinean point —the methodology of empirical belief revision

Can it be assumed that anti-exceptionalists, as good card-carrying analytically minded philosophers, will be hospitable to the idea that the rational revision of systems of empirical belief will be subject to a *principled methodology*— an explicit set of precepts for qualifying what should and should not be accepted in the light of incoming evidence? If this is a safe assumption, what can be said about the logical principles and inferential practices that will be called upon in explicating any specific application of such precepts?

The reader will doubtless anticipate what's coming. If a Quinean empiricist's revision procedures allow of systematic principled articulation, the methodological principles concerned will presumably assume the kind of general conditional form schematised in *Rules-in-General* earlier, namely,

[37] The reader should reflect that versions of instantiation and detachment will inevitably feature in any adequate 'logic of proof-checking'.

> If such-and-such a situation arises, consider dropping that principle/hypothesis that meets such-and-such conditions and replacing it with...[38]

So in order to articulate the explicit *rationale* for a particular revisionary proposal that is sanctioned by such a principle, a theorist will need to advert to the mandate for it that is entailed by the methodological precept(s) concerned, in conjunction with some description of the specific predicament into which the theory in question has fallen. And now it seems unintelligible how that articulation could avoid involving reasoning by instantiation and detachment. In addition, since it seems that spelling out the idea of 'recalcitrance' of evidence for theory will need appeal to a notion of inconsistency, an explicit justification for the idea that revision is even called for in the first place will need to appeal to the validity of some version of a negation-introduction rule.

In sum: the rationalisation of revisions made to logic in accordance with a broadly Quinean methodology, or indeed with any methodology, must, we contend, advert to the presumed validity of entailments sanctioned by some combination of something like *MPP* and *UI*. Again, this is consistent with allowing 'tweaks' to or restrictions on these principles. But some versions of conditional and instantiational thought, and some form of principled aversion and inferential response to inconsistency between theory and data, are quintessential to anything recognisable as a coherent methodology of theory revision at all, including principles of revision of logic. Whatever exact shape the relevant versions of *UI* and *MPP* take for particular theorists, the general principles they exemplify will need to maintain a position above the Quinean fray.[39]

References

Berger, Alan (2011) *"Kripke on the Incoherency of Adopting a Logic"*, Ch. 8 of his Saul Kripke. Cambridge: Cambridge University Press.

[38] This formulation is, of course, extremely crude. It would be for a Quinean theorist to give more exact thought to a formulation of the principles in their preferred methodology of empirical belief addition and revision. But the important point is that, whatever the theorist's preferred formulation of such rules, their formulation will be general and their application to particular situations will require instantiation and detachment (in whatever specific forms).

[39] For valuable discussion, we would like to thank the participants at the Kripke/Birman virtual seminar on the 'Adoption Problem' in the Fall of 2020, especially Saul Kripke and Romina Birman, the members of the Stirling/Templeton *Knowledge Beyond Natural Science* A Priori group, and the participants at the Rome Conference on Normativity in the summer of 2022. Two anonymous referees for *MIND* provided invaluable critical comments.

Besson, Corine "Knowledge of Logical Generality and the Possibility of Deductive Reasoning", forthcoming in Chang and Nes (eds), *Inference and Consciousness*. London: Routledge.
Birman, Romina (formerly Romina Padro) (2023) "[title]", *Mind*, [vol] ([issue]), [page range].
Boghossian, Paul (ms) "Guidance by Rules".
Boghossian, Paul (2003) "Blind reasoning", *Aristotelian Society Supplementary Volume* 77 (1), pp. 225-248.
Boghossian, Paul (2008) "Epistemic Rules", *The Journal of Philosophy*, 105(9), pp. 472-500.
Boghossian, Paul (2017) "Postscript: Further Thoughts about Analyticity, 20 Years Later", in Wright, Hale and Miller (eds), *A Companion to the Philosophy of Language*, Second Edition. Blackwell, pp. 611-618.
Boghossian, Paul and Timothy Williamson (2020) *Debating the A Priori.* Oxford: Oxford University Press.
Casalegno, Paolo (2004) "Logical Concepts and Logical Inferences", *Dialectica*, 58, pp. 395-411.
Devitt, Michael and Jillian Rose Roberts (2023) "Changing our Logic: a Quinean Picture," *Mind*, [vol]([issue]), [page range].
Field, Hartry (2005) "Recent Debates about the A Priori", in Gendler and Hawthorne, (eds), *Oxford Studies in Epistemology.* Oxford: Oxford University Press, pp. 69-88.
Finn, Suki (2019) "Limiting Logical Pluralism", *Synthese* online first, doi:10.1007/s11229-019-02134-8.
Hale, Bob and Crispin Wright (2000) "Implicit Definition and the A Priori", in Boghossian & Peacocke (eds), *New Essays on the a Priori*. Oxford University Press, pp. 286-319
Harman, Gilbert (1986) *"Change in View: Principles of Reasoning"*. Cambridge: MIT Press.
Kripke, Saul (2023) "The Question of Logic", *Mind*, [vol]([issue]), [page range].
Maddy, Penelope (2002) "A naturalistic look at logic", *Proceedings and addresses of the American Philosophical Association*, 76(2), pp. 61-90.
Padro, Romina (now Romina Birman) (2015) unpublished CUNY PhD thesis, "What the Tortoise Said to Kripke: the Adoption Problem and the Epistemology of Logic", downloadable at https://academicworks.cuny.edu/gc_etds/603/
Priest, Graham (2014) "Revising logic", in Rush (ed.), *The metaphysics of logic, chapter 12*. Cambridge: CUP, pp. 211-223.

Putnam, Hilary (1969) "Is Logic Empirical?", in *Boston Studies in the Philosophy of Science*, vol. V, pp. 216-241.

Quine, W.V.O. (1951) "Two Dogmas of Empiricism", in *The Philosophical Review*, 60 (1), pp. 20-43. doi:10.2307/2181906. Reprinted in his (1953) *From a Logical Point of View*, Harvard University Press.

Williamson, Timothy (2007) *The philosophy of philosophy.* Oxford: Blackwell.

Williamson, Timothy (2013) *Modal logic as metaphysics*. Oxford: Oxford University Press.

