

Against Reflective Equilibrium for Logical Theorizing*

Forthcoming in *Australasian Journal of Logic*

Jack Woods

September 13, 2018

Anti-exceptionalism about logic denies that our logical theories and our choice thereof have any special epistemic character. In particular, however we should go about choosing a logical theory, it will be analogous to how we choose other non-logical theories. This view originates with Quine (1970) and Goodman (1983) and has become increasingly common as a picture of the status of logic.¹ It promises to resolve deep issues about how to discriminate between alternative options for our background logic; that is, our most general canons of implication. But promise is one thing, ability another; I'll suggest here some limitations of this approach to logical revision.

I'll start by distinguishing two ways of developing anti-exceptionalist approaches to logical revision. The first emphasizes comparing the theoretical virtuousness of developed bodies of logical theories, such as classical and intuitionistic logic. I'll call this *whole theory comparison*. The second attempts local repairs to problematic bits of our logical theories, such as dropping excluded middle (and modifying elsewhere accordingly) to deal with intuitions about vagueness. I'll call this the *piecemeal approach*.

I'll then briefly discuss a problem I've developed elsewhere for comparisons of logical theories. Essentially, the problem is that a pair of logics may each evaluate the alternative as superior to themselves, resulting in oscillation between logical options. The piecemeal approach offers an way out of this problem and thereby might seem a preferable to whole theory comparisons. I'll go on to

*Thanks to Liam Kofi Bright, Edd Elliot, Daniel Elstein, Ole Hjortland, Beau Madison Mount, Erum Naqvi, Carlo Nicolai, Graham Priest, Gillian Russell, Gil Sagi, Dan Waxman, Robbie Williams, and Timothy Williamson for helpful discussion of these issues. Thanks to the Centre for Metaphysics and Mind at Leeds for providing useful feedback for an earlier version. And thanks to my comprehensive exam examiners—John Burgess, Tom Kelly, Boris Kment, and the sadly departed Delia Graff Fara—who heard an *extremely* early version of this work many years ago.

¹Recent sophisticated articulations include Otavio Bueno and Mark Colyvan (2004), Ole Hjortland (2017), Graham Priest (2006, 2016), Michael Resnik (1997), Gillian Russell (2013, 2015), and Timothy Williamson (2014, 2015).

show that reflective equilibrium, the best known piecemeal method, has deep problems of its own when applied to logic (by developing a problem suggested by (Wright 1986).)

The problem developed isn't unique to reflective equilibrium, but affects any piecemeal method which both (a) allows significant freedom in how to revise and (b) *uses* a logical theory to assess the reasonableness of various repairs. (a) and (b) together yield a far too permissive account of logical revision. Such approaches thereby require some kind of *guidance* to separate out good repairs from bad ones. However, it's unclear how to add this guidance in an epistemically satisfactory way without abandoning the anti-exceptionalist standpoint.

I'll close by suggesting that the best version of the piecemeal method of revision makes use of multiple theoretical virtues and, moreover, adopts a constraint—LOGICAL PARTISANHOOD—which was originally motivated for whole theory comparisons (Woods forthcoming-a). The upshot of my discussion is that

- piecemeal revision of our logical theories, especially in its reflective equilibrium guise, suffers from distinctive problems,
- solving these problems requires epistemically significant guidance, but it's difficult to do so without abandoning anti-exceptionalism,
- the best way resolving these problems requires significantly constraint, so much so that it's not really recognizable as a form of reflective equilibrium.

The argument and discussion that follows will be kept slightly informal to avoid complication and theoretical commitments not essential to the overall point.

0.1 RULES OF THE GAME

A few reasonable presumptions frame my discussion. First, I assume that anti-exceptionalist views of logical theory choice are really about which logic to adopt as our most basic canon of legitimate deductive implication. This is because deciding which particular logics to adopt for instrumental purposes poses no interesting philosophical difficulties. We can always treat these advantages as features of using a particular formal method, like relevance logic, in a particular context *without* abandoning the governing role of our (roughly classical) background canons of implication.²

Certain familiar arguments mustered in favor of logical deviance, like the usefulness of linear logic in tracking resource use or the usefulness of relevance logic in theorizing informational content, are thereby shown to be undercooked.³ We

²And, as I've argued (forthcoming-a), certain nice advantages of strongly deviant logic can, as yet, *only* be seen from the viewpoint of our background canons of implication.

³See Shapiro (2014) for a useful survey of applications of this type. Shapiro himself questions the coherence of a single background canon of deductive implication, but there's no space to engage with his arguments for that conclusion here.

need an additional reason to think that the explanation just offered—in terms of situation-specific reasoning—isn't sufficient to account for the supposed advantage. I doubt that such a reason will be forthcoming, but regardless there remain interesting questions about logical revision which focus explicitly on our most general canons of implication. It's these questions, like which logic we should take as our background logic in which to evaluate logical relations between propositions, that I focus on.

Second, anti-exceptionalist methodology aims to provide us with the ability to justify this or that logical revision. In carrying it out we thereby need to look for reasons we're in a position to use in providing justifications. Procedures of idealization which are available in other areas aren't clearly available since what counts as an idealization *itself* involves implication (see §0.4.1). Similarly, evidentialist ways of cashing out notions of epistemic reasons also take for granted a fixed background logic since they employ notions like 'independent evidence' which are explicated using an implication relation. There simply isn't a theory-neutral stance from which to make sense of reasons which exist but which we're not able to make use of in justifying a case for revision.

This situation places important “dialectical”, “internalist”, “basing”-looking constraints on the methodology that become important below. Reasons which can be generated from some logical resources distinct from our own are thus unavailable (see §0.4.1 and (Woods forthcoming-a) for more detailed discussion.)⁴

0.2 TWO METHODOLOGIES FOR LOGICAL REVISION

Anti-exceptionalist methodology comes in two rough forms. First, some treat evaluation of our logical options much like we treat evaluation of whole scientific theories more generally. This involves evaluating whole bodies of developed logical theory according to various criteria of theoretical goodness. Call this *whole-theory comparison*. Others treat evaluation of our logical options more *piecemeal*, employing methodologies like wide reflective equilibrium (Goodman 1953) in developing and justifying our logical commitments.

The advantages of whole-theory comparison in actual use is that we know quite a bit about the various alternatives before comparing them. To be even eligible for comparison in this way, a certain amount of preparatory work needs to be done. Moreover, typical alternatives are more or less natural and guided by some background idea. This means that the packages tend not to be logical frankensteins. There are, of course, exceptions.

The advantage of the piecemeal methodology in actual use is its ability to innovate, determining by means of a series of tweaks an overall theory cut to fit

⁴I write in this slightly elliptical way since “possessing” beliefs or knowledge that would suffice for justification is tricky when our fundamental logical resources are up for grabs. What I have in mind should nevertheless be clear.

the data. Yet it works more blindly since the consequences of various tweaks are often not at all obvious. Moreover, the resulting series of tweaks can look quite frankensteinian.

Of course, the methods are related and, in fact, from a certain remove one can see the piecemeal approach as a sequence of whole-theory comparisons of underdeveloped theories. Nonetheless, it's important to separate them here. Separating them makes it clearer that piecemeal methodology—at least when explicated in certain ways—has even more problems than whole-theory comparisons.

Why might we be tempted by the piecemeal approach? One reason is because of standing problems with whole-theory comparisons. I've explored these problems elsewhere, arguing that whole-theory comparison needs to be significantly constrained in order to play the role many have suggested it should (Woods forthcoming-a). Let me explain.

Evaluation of our logical alternatives, in terms of their theoretical goodness, makes use of logic. We should use the logic we currently endorse—our *working logic*—in order to evaluate the benefits of alternatives.⁵ But our working logic sometimes finds merits of an evaluated logic which aren't available from its own perspective; this is especially the case when an evaluated logic is so weak that developing a sufficiently strong syntax and metatheory is extremely difficult.⁶ So our evaluation of the theoretical virtues can be significantly biased by which logic we currently accept.

If, then, we adopt the usual simple choice principle—move to an alternative logic when it does better than our own according to a privileged set of theoretical virtues—then we can be forced to oscillate between a pair of logical alternatives as each looks better than the other by their opponent's lights. Call this the *oscillation problem*.

This problem can be solved by resolving only to adopt a logical theory which is marked not inferior by our current logic if it is also marked not inferior by its own lights.⁷ In practice, this constraint favors cases of whole theory comparisons where there's sufficient overlap between the compared logics in which to do basic metatheory, including a theory of syntax. The solution rules out many otherwise interesting logical disputes since the metatheory of many logical alternatives is seriously underdeveloped at best.

⁵The situation is more complicated than this, but details would distract here. See both (Woods forthcoming-a) and below.

⁶Many recent weak non-classical logics have this problem, as pointed out by Meadows and Weber (2016) and myself (forthcoming-a).

⁷Why 'not inferior' instead of 'superior'? This isn't perversity, it's to allow that in cases of ties, we *can*, though we need not, switch to an alternative. Let no one call me unecumenical.

It's thus natural to look for different ways of investigating logical alternatives that don't effectively rule out seriously weak logics in advance. And there's a natural candidate: perhaps we should focus on piecemeal revisions of our existing logical theories instead of engaging in comparisons of whole developed bodies of logical theory. This has the potential to avoid the worries and problems, like the one we've just discussed, which arise from using our logic in the course of revising our logic. For instance, Hjortland writes:

An abductive argument for a logical theory will inevitably presuppose some laws of logic, but that is not incompatible with revision of logic. All the laws of logic cannot be subject to revision simultaneously, nor is that a requirement. The anti-exceptionalist only needs to hold that no law of logic will be beyond revision. (2016)⁸

I take it that the suggestion here is that we should evaluate our logical theories from *within*, making changes so as to improve their overall coherence and resolve problems. This in turn suggests using a method like reflective equilibrium in revising logic and, in fact, reflective equilibrium avoids the most straightforward versions of the oscillation problem.⁹ Reflective equilibrium is our most developed and widely accepted general method for piecemeal revision of theories, so it's the natural candidate for logical revision as well. At least *prima facie*.

More complicated pictures of piecemeal logical revision are also possible, but they plausibly fall victim to similar, if fancied up, version of the complaints I'll raise below.¹⁰ Reflective equilibrium will thereby serve as our well-worked out and probative example of problems for piecemeal approaches to logical revision. This is because of two things: first, narrowing the relevant virtues to something like coherence, as in reflective equilibrium, can easily be seen to avoid the problem of some virtues pointing one way, other virtues another. Second and relatedly, there is tremendous flexibility in *how* to revise our logic so as to increase coherence. Given these two facts, reflective equilibrium looks more promising as a methodology for anti-exceptionalist accounts of logic. But again, as always, promise is one thing and ability another. We now turn to some details.

0.3 REFLECTIVE EQUILIBRIUM AND LOGICAL REVISION

Reflective equilibrium was first explicitly sketched by Goodman (1953). Interestingly, his example of a place where we obviously use this method was logic, though it took over half a century before people explicitly followed him in this.¹¹ The method of reflective equilibrium presumes a distinction between two sorts

⁸See also Russell (2014), though she is more charitably read as providing an etiological story about our logical beliefs instead of a justificatory one.

⁹Hjortland doesn't explicitly endorse this particular methodology, but it's a natural explanation of the underlying thought.

¹⁰See Mares (2014) for an example of an exquisitely spelled out piecemeal approach.

¹¹I say "interestingly" because one would have expected Goodman to pick up on the problem I'll lay out. Perhaps, though, Goodman was simply too influenced by Quine at the time. This wouldn't be surprising given that Quine's account of logic—in terms of the "web of belief"—is

of “beliefs” about some subject matter.¹² The core of reflective equilibrium is a methodology for bringing these two set of belief into harmony.

One set of beliefs comprises judgments and intuitions about particular cases of implication.¹³ These dispositions to judge that particular statements imply other statements and beliefs that such and so particular immediate inferences are legitimate are our raw data—they function somewhat similarly to observations and observation statements. We can divide this raw data up into two sets— \mathfrak{L}_p , those dispositional beliefs we take to be probative and \mathfrak{L}_{np} , those we don’t (writing the pair of these as $\mathfrak{L}_{p/np}$.) We shift material between these two sets in the course of our investigation.

We’re disposed to recognize that Socrates is wise and good implies that Socrates is wise and the corresponding inference as entirely unobjectionable. It’s a clear member of \mathfrak{L}_p . We’re also disposed to believe

$$\exists\varphi, \psi \neg(\varphi \rightarrow \psi) \wedge \neg(\psi \rightarrow \varphi)$$

is true; this belief is very plausibly an initial member of \mathfrak{L}_p . When we come to treat \rightarrow as the material implication, those of us who are sensible move our dispositional belief in the truth of $\exists\varphi, \psi \neg(\varphi \rightarrow \psi) \wedge \neg(\psi \rightarrow \varphi)$ from \mathfrak{L}_p to \mathfrak{L}_{np} . Since we’ve a justification of the correctness of $\forall\varphi, \psi (\varphi \rightarrow \psi) \vee (\psi \rightarrow \varphi)$ on the basis of other principles we accept, we come to see our initial intuition as non-probative.

The other set of beliefs collects up our views about generally valid forms of implication and inference.¹⁴ This set, \mathfrak{L}_G , contains such banalities as the belief that conjunctions imply their conjuncts and that inferring a conjunct from a conjunction is universally legitimate, as well as more technical beliefs like the quantifier-negation commutation rules and contentious beliefs like the (logical) validity of the ω -rule.

This might suggest that a logic is simply a set of (logically true) statements, a picture strongly associated with Quine’s account of logic. This would be problematic—as has been repeatedly pointed out over the years, this aspect of Quine’s picture is simply inadequate in theorizing different logics. Quite different logics generate the same set of logical truths and some quite common logics, like *K3*, have no logical truths at all. Reflective equilibrium is not at all committed to this mistake; rather, our sets of general beliefs will include

strikingly similar to Goodman’s. Here I treat both as instances of the reflective equilibrium picture, flagging where the differences matter.

¹²I’ll speak loosely here, lumping all kinds of representational states together as beliefs. Pedantic precision would distract here.

¹³Of course, we also want to distinguish between implication and inference more generally, but our intuitions aren’t as clean as those who make this distinction would like. I’ll not fuss much about the distinction between these since it’s generally orthogonal to the points I want to make.

¹⁴Of course, distinguishing between \mathfrak{L}_G and $\mathfrak{L}_{p/np}$ in actual practice isn’t trivial, not by a long shot. As my aim here is primarily critical, I won’t worry about this.

all kinds of claims about which argument forms are valid and which inferential transitions are legitimate.

The content of the beliefs in \mathcal{L}_G predicts the validity of certain particular cases of implication—we’ll write this prediction relationship $\mathcal{L}_G \Rightarrow \mathcal{J}$. In the best case, the predictions of \mathcal{L}_G won’t clash with \mathcal{L}_p ; when they do—when there’s $q \in \mathcal{L}_p$ such that q claims \mathcal{J} as incorrect—we have tension (writing this as $\perp(\mathcal{L}_G, q, \mathcal{J})$.) There can be relative degrees of *coherence* between \mathcal{L}_G and \mathcal{L}_p , depending on the number and “depth” of instances of \perp .¹⁵ Tension, of various degrees, is common even in logic. Since many such tensions don’t wear their solution on their face and since investigation over the years shows costs and benefits of any solution, this fact itself motivates some versions of anti-exceptionalism.

For example, while most theorists think that *modus ponens* is generally valid, some of these theorists also have the intuition that McGee-style instances (1985) aren’t. Personally, I think the introduction and elimination laws for the truth-predicate are logically valid, as are both *modus ponens* and conditional proof, even though English is semantically closed.¹⁶ However, I don’t *want* to recognize the triviality that follows from these four beliefs by means of Curry’s paradox (or, god forbid, reject a structural rule). And so on.

To improve the coherence of our collection of beliefs and, thereby, to improve our justification for maintaining them, we tweak our system of beliefs in a systematic fashion. First we apply the rules endorsed by \mathcal{L}_G in order to generate predictions which cause conflict. In order to resolve the conflict, we revise \mathcal{L}_G to fit with \mathcal{L}_p and, correspondingly, move members of \mathcal{L}_p to \mathcal{L}_{np} (and conversely) to reduce tension with \mathcal{L}_G . Then we reapply the rules again to generate more conflict (or, in the best case, to delightfully see that all apparent conflict has been resolved.)

The aim is to improve the coherence between the predictions of the general beliefs and the “observational” data of our dispositional beliefs about particular cases. If we take coherence of these two to be *a* or, even more strongly, *the* significant good-making epistemic feature of a set of beliefs, then through this process the resulting package of beliefs about logic becomes progressively more epistemically justified.¹⁷

When the whole body of such beliefs reaches a point where disagreement

¹⁵We need something like depth to avoid simply measuring coherence by counting instances of tension; intuitively, some tension is more dramatic than others. Again, details would distract from the main point so I’ll leave this at an intuitive level. See §0.5.2 for a strategy on which “depth” is taken more seriously.

¹⁶Or, alternatively, my arithmetical beliefs legitimate arithmetical resources extending Robinson arithmetic. In my case, both.

¹⁷See McPherson (2015) for trenchant arguments that reflective equilibrium falters here.

has been maximally minimized—for example, when all dispositions to judge predicted implications as invalid have ceased, been explained away, or cease to be predicted—then the system is in *reflective equilibrium*.¹⁸ When this is just coherence between \mathfrak{L}_G and $\mathfrak{L}_{p/np}$, then we say that the theory is in *narrow* reflective equilibrium.

In the case of logic, we need other facts besides those just concerning implication in order to explain away our intuitions. We need psychology to explain away the widespread tendency for people to fail to recognize cases of modus tollens. Likewise, we should probably demand that our logical theories cohere with background constraints such as the ability to recapture obviously true contentual mathematics like elementary fragments of arithmetic.¹⁹

When our logical theory and intuitions achieves both internal coherence and coherence with a large body of presupposed background theory from other areas, we say it is in *wide* reflective equilibrium. It seems we should aim for at least *somewhat* wide reflective equilibrium, given the above examples. I’ll assume so moving forward.²⁰

This brief gloss on reflective equilibrium will do for our purposes. We won’t make much hay of the distinction between wide or narrow reflective equilibrium or focus overly much on the endpoint of the process. Our focus will be on the process itself and on whether successive revisions confers a form of epistemic justification on the theory so arrived at.

0.3.1 EVALUATING COHERENCE AND PREDICTION

Observe that we didn’t define ‘coherence’, ‘prediction’, or ‘confirmation’. This was for good reason: the best accounts of these notions—in fact, all reasonable accounts of these, good or not—make essential use of an implication relation. This means that we need to settle which account of implication will govern how we evaluate the coherence of our logical theory, what our general beliefs predict, and how our dispositional beliefs confirm or not our general beliefs. This is a problem local to the application of reflective equilibrium to logic and perhaps elementary fragments of mathematics. Applying this methodology to most non-logical matters, this issue doesn’t arise as the relevant theoretical resources involved in explicating coherence, prediction, and confirmation are held fixed.²¹

¹⁸Note that the property of being in reflective equilibrium is really defined internally to \mathfrak{L}_G . In principle, this means that it’s possible that according to one way of cashing out what \mathfrak{L}_G predicts, our beliefs are in reflective equilibrium, but according to another way, it’s not. This is especially pertinent when dealing with disputes about logic; see below.

¹⁹We need the ability to theorize about notions of proof and derivability as well (see §0.4.1), but it’s not clear whether or not to treat this as part of \mathfrak{L}_G . I’ll assume so going forward, for simplicity, flagging up where it matters.

²⁰For more examples which could be massaged into showing that narrow equilibrium isn’t sufficient, see Williamson (manuscript).

²¹Another potential exception is theoretical normativity itself, though it’s worth noting that the problem appears most drastic in the case of logic. See Woods (2018) for discussion.

When comparing a pair of logics, we have a few options about which logic to use in assessing their various merits—in particular, in drawing out the extensions of \Rightarrow and \perp .²² We could use our current logic—here something like what’s articulated by \mathfrak{L}_G —in assessing the coherence and predictions of \mathfrak{L}_G and an alternative $\mathfrak{L}_{G'}$. Alternatively, we could use \mathfrak{L}_G to assess the coherence and predictions of \mathfrak{L}_G and $\mathfrak{L}_{G'}$ to assess the coherence and predictions of $\mathfrak{L}_{G'}$. Both are natural ways of proceeding, at least at first glance.

But, on second glance, it’s significantly more reasonable to use \mathfrak{L}_G to assess the predictions and coherence of \mathfrak{L}_G and $\mathfrak{L}_{p/np}$ and $\mathfrak{L}_{G'}$ to assess the predictions and coherence of $\mathfrak{L}_{G'}$ and $\mathfrak{L}_{p'/np'}$. What about using \mathfrak{L}_G , our starting set of logical beliefs, to assess the coherence of \mathfrak{L}_G and $\mathfrak{L}_{p/np}$ as well as the coherence of $\mathfrak{L}_{G'}$ and $\mathfrak{L}_{p'/np'}$? That is, what about using our current logical beliefs in order to assess what both our current logic predicts and what the effects of resolving a bit of incoherence would be? This, turns out, is unreasonable. Many problems due to members of \mathfrak{L}_G will persist, by the lights of \mathfrak{L}_G , even after we move from \mathfrak{L}_G to $\mathfrak{L}_{G'}$.

Suppose $\mathfrak{L}_G \Rightarrow \Delta \models \varphi$, for Δ “in” \mathfrak{L}_G and where \mathfrak{L}_p holds that φ isn’t correct—i.e. where we have $\perp(\mathfrak{L}_G, p, \neg\varphi)$. One natural way to revise our beliefs is to ensure that $\Delta \models \varphi$ isn’t predicted by $\mathfrak{L}_{G'}$, by making corresponding changes to the beliefs supporting this entailment in \mathfrak{L}_G . If, however, we evaluate the predictions of $\mathfrak{L}_{G'}$ by \mathfrak{L}_G , presuming even a minimal amount of closure, we’ve got a problem. Δ is still in \mathfrak{L}_G , so $\mathfrak{L}_G \Rightarrow \varphi$, and so we have $\perp(\mathfrak{L}_{G'}, p, \neg\varphi)$ and our tension remains unresolved.

Since reflective equilibrium targets a bit of tension then resolves it, then the above situation will tempt us to move our intuition about φ from \mathfrak{L}_p to \mathfrak{L}_{np} . After all, this really does resolve the tension and Δ might be more or less inviolate. But continually resolving $\mathfrak{L}_{p/np}$ —instead of sometimes revising what \mathfrak{L}_G predicts—seems a bad tendency in applying the method of reflective equilibrium, undermining much of its intuitive appeal. \mathfrak{L}_G and $\mathfrak{L}_{p/np}$ are supposed to be roughly equal players in the game of reflective equilibrium; this fact seems essential to reflective equilibrium conveying epistemic justification on its outcome.

So we should be less ecumenical about which logic to use in coming up with our revisions here; for simple applications of reflective equilibrium, we should use \mathfrak{L}_G to assess its own predictions and tension with $\mathfrak{L}_{p/np}$ for all stages of \mathfrak{L}_G and $\mathfrak{L}_{p/np}$. That is, as we revise our logical beliefs, we should continually update our account of what they predict in line with these revisions; moreover, when considering a potential revision, we should assess what would result by the lights of the proposed revision, not by the lights of our current logical beliefs. We can (and will) finesse the problem for now by treating the extension of \Rightarrow

²²See Woods (forthcoming-a) for further discussion.

and \perp as themselves parts of \mathfrak{L}_G . This means that \perp and \Rightarrow are generated ‘pointwise’ at every stage of revision.²³ We’ll relax this restriction below in the positive proposal section.

0.3.2 REFLECTIVE EQUILIBRIUM V. WHOLE THEORY COMPARISON

Again, the differences between the picture I’ve just spelled out and the whole theory comparison picture may seem mostly ones of emphasis and application. And, again, reflective equilibrium can be seen as whole theory comparisons with a single theoretical virtue—coherence between \mathfrak{L}_G and $\mathfrak{L}_{p/np}$ (and our broader scientific beliefs)—and a narrow scope of available alternatives. There will hardly be the body of developed logical theory that’s characteristic of the logical alternatives—intuitionistic, relevantistic, etc.—typically being compared in whole logic comparisons.

When comparing the costs of two developed bodies of logical theory, like intuitionistic and classical logic, we have decades of refinement at our disposal and a determinate sense of what can and cannot be done in each. But as the majority of proposed revisions we encounter in carrying out reflective equilibrium will be novel (and many extremely surprising), we have significantly less knowledge of the consequences of revising while so revising. *A fortiori*, we have significantly less knowledge to work with in building a case for revising any particular way. So even though we *could* treat reflective equilibrium like a case of whole-theory comparison, this would obfuscate potentially epistemically significant differences that our pre-existing knowledge makes to the reasons we actually have when evaluating logics.

More interestingly for our purposes here, it’s the restriction to one governing theoretical virtue—coherence—that is useful in dealing with our initial motivating problem. Since we are only evaluating coherence, it’s more difficult to find cases where successive revisions result in cycles. That would require that small coherence-improving revisions could land you back into the same place over time; this is possible, but it’s difficult to come up with an especially compelling determinate case for it. The restriction to a single theoretical virtue makes it easier to avoid cases where cross-cutting theoretical virtues will push us back and forth in the course of successive revisions.

We’ll thus assume for present purposes that simple reflective equilibrium improves over typical cases of whole-theory comparisons with respect to worries about revision cycles. So presumed, reflective equilibrium yields a fairly straightforward and compelling account of how our logical beliefs are justified. It doesn’t require any form of a priori knowledge or magical capacity of rational insight to make sense of how our beliefs are justified. It also fits nicely with

²³Of course, the problems raised below can be iterated given this assumption, but let that pass. The situation is bad enough as it is.

historical facts about the piecemeal development of contemporary logical orthodoxy over time, such as our gradual relinquishing of existential import.²⁴

And it puts the justification of logical beliefs on a par with widely endorsed views about the proper methodology in developing a naturalistic account of other non-empirical matters—inductive methods (Goodman 1953), justice (Rawls 1971), normativity (Scanlon 2014), and so on. However, the case of logic is different from these since logic occupies two roles in reflective equilibrium—as both the object of our theorizing and as the arbiter of our theorizing. When the target of reflective equilibrium occupies these two roles an interesting problem emerges.

0.4 THE DEGREES OF FREEDOM PROBLEM — INTUITIVE VERSION

This problem was first raised during Crispin Wright’s criticism of Quine’s views on logic (1986). It was subsequently refined by Stewart Shapiro in his criticism of Resnik’s (1997):

Suppose that a logician has an intuition that a certain argument A is invalid, and wants to see if this intuition coheres with her evolving logical theory T . Sadly, she finds out that the invalidity conflicts with T . Consider the sentence:

(*) The theory T is not in accord with the invalidity of the Argument A

presumably accepted by the theorist. We are told that any sentence is up for revision. Can our logician maintain both T and the invalidity of A by rejecting (*)? That is, can our logician just reject the inference from T to the validity of A ? Regress threatens. (2000)

Regress maybe, but the real issue here seems different to me. It’s that we simply have too many options—we have too many *degrees of freedom*—in how to increase the coherence of our logical theories in the face of recalcitrance. I’ll start by raising the problem in an intuitive fashion—i.e. without further analyzing prediction or conflict—then briefly discuss more complicated reconstructions.

Suppose we have a dispositional belief p that a certain implication \mathfrak{J} is invalid, yet our general beliefs predict its validity. We have two *prima facie* legitimate ways of revising our logical beliefs: (1) revise the bit of \mathfrak{L}_G that’s predicting \mathfrak{J} and (2) move p from \mathfrak{L}_p to \mathfrak{L}_{np} . And, in both cases, making other reasonable adjustments like dropping conjunctions in these sets which contain the revised bit, etc.

²⁴See Church (1964) for a history of existential import and Hanson (1989) for a compelling case for its status as a recalcitrance in our logical theorizing.

However, reflective equilibrium also permits the following moves: (3) remove $\mathcal{L}_G \Rightarrow \mathcal{J}$ from \mathcal{L}_G , leaving much of the rest alone, and (4) remove $\perp(p, \mathcal{L}_G, \mathcal{J})$ from \mathcal{L}_G , leaving much of the rest alone. After all, these two cited facts are, by assumption (though see below), just further members of \mathcal{L}_G . On the Quinoid web-of-belief version of the anti-exceptionalist picture that Shapiro and Wright target, it's explicit that $\mathcal{L}_G \Rightarrow \mathcal{J}$ and $\perp(p, \mathcal{L}_G, \mathcal{J})$ are just further bits of theory. If they're just more beliefs in \mathcal{L}_G , then denying their revisability fact treats the logic of evaluation as special. This is strongly at odds with anti-exceptionalism. So the anti-exceptionalist needs to allow that facts about \perp and \Rightarrow can also be revised.

Both (3) and (4) remove the tension with $\mathcal{L}_{p/np}$. The tension is represented by $\perp(p, \mathcal{L}_G, \mathcal{J})$ and predicted by $\mathcal{L}_G \Rightarrow \mathcal{J}$ so removing either alleviates the incoherence. In doing so, we cease to believe that \mathcal{L}_G predicted tension or cease to believe that the predictions of \mathcal{L}_G were tension-involving. Given that these revisions leave our intuitions untouched and preserve the central portions of \mathcal{L}_G , they are *officially* minimally damaging and thereby attractive revisions.

Unofficially and actually, they're grotesque. We should not be able to simply revise away tension by removing our belief in it—after all, it's still there—or denying that our general logical beliefs predict it. Revising logic in this way is rather like hiring a plumber to fix a leaky pipe, then watching them look at the growing puddle of water on the floor and insist the pipe is fine. Even if we believe them, we're gonna have soggy carpet.

It seems of rather little use to have a “coherent” system of beliefs when whenever we were faced with conflict we could *always* modify what we believe \mathcal{L}_G predicts so as resolve the “incoherence”. And system of beliefs in reflective equilibrium which results from a sequence of such moves will be more or less useless, at least by our current lights. And if we can't say that from within the terminal state of our process of reflective equilibrium, so much the worse for reflective equilibrium.

Sometimes modifying \Rightarrow or \perp will be the correct move, but it can't always be both available and attractive. It seems ridiculous that when we have our theory's predictions disconfirmed, we can always recover by simply denying that there is any disconfirmation. The process of reflective equilibrium is supposed to be a process by which we match theory to evidence and evidence to theory. It's pointless to engage in such a project if we can get ourselves into reflective equilibrium by simply revising away our beliefs about coherence or prediction.

0.4.1 COMPLICATING THE DEGREES OF FREEDOM PROBLEM?

The version of degrees of freedom problem just given presumed that coherence and prediction were to be represented by the beliefs in \mathcal{L}_G which involve \perp and \Rightarrow . Presumably these are somewhat informed by the other members of \mathcal{L}_G as

well, but essentially what's predicated by \mathcal{L}_G is what \mathcal{L}_G more or less directly says is predicated by \mathcal{L}_G .

We might try to avoid this problem by explicitly defining coherence and prediction. For instance, we might try defining \Rightarrow as follows: $\mathcal{L}_G \Rightarrow \varphi$ just in case φ can be derived by repeated application of the “rules” endorsed by \mathcal{L}_G .²⁵ Likewise, we could treat coherence in terms of \mathcal{L}_G predicting or containing some claim φ whose contradictory was in \mathcal{L}_p .

Unfortunately, this doesn't move us far enough from the original version of the problem. Since reflective equilibrium is supposed to be a method we can actually employ, conveying justifications we can cite in favor of our revisions, medical limitations restrict our ability to use coherence and prediction in revising our logical theories. This is one of those important ‘basing’-type restrictions we mentioned above. We need to be able to recognize and display that our move is justified.

We thus need to use what we (defeasibly) *justifiably believe* about prediction and coherence instead of what actually holds about prediction and coherence in revising. In practice, this means that we need to have a minimal theory of derivability in order to have a suitably expansive collection of beliefs about prediction and coherence. This in turn means that we need to have a suitably strong theory of syntax and inductive definitions in order to carry out the basic metatheoretic results which undergird our beliefs about what \mathcal{L}_G predicts. And *these* are certainly part of \mathcal{L}_G and thereby revisable on the anti-exceptionalist picture.²⁶

If our beliefs about this metatheory are themselves revisable—if they aren't privileged over other members of \mathcal{L}_G —then we have not improved over our intuitive version of the problem. We can still revise away any seeming tension or bad prediction by revising—presumably in somewhat ad hoc ways—our theories of syntax and inductive definitions. Of course, ad hoc revisions are bad in some way, but our methodology does not by itself make any use of theoretical virtues beyond coherence in permissible revisions; nor does wide reflective equilibrium itself require our beliefs to be non-gerrymandered.²⁷

It's tempting to lean back on some kind of idealization to avoid this problem

²⁵ Again, as above, we're leaving this informal. Strictly speaking we'd need to somehow isolate the rules from other members of \mathcal{L}_G .

²⁶ Alternatively, we could treat these as part of the widened body of beliefs that \mathcal{L}_G needs to cohere with. Even then, we can modify their impact by modifying instances of \Rightarrow involving these materials. Nothing substantial turns on this, so I'll presume they're part of \mathcal{L}_G going forward.

²⁷ Insofar as we're willing to block ad hoc revisions, we could have already used the typical ad-hoc-ness of (3) and (4) to focus on the two reasonable options (1) and (2) for proceeding in the face of tension. And do so without giving the more detailed story about prediction and coherence. See below.

and, in other contexts, this is unproblematic. So long as enough of our theoretical tools are “fixed” for the purposes of application of reflective equilibrium, then we can use the idealized versions of \Rightarrow and \perp to avoid the degrees of freedom problem. But when logic is the target of our investigation, then exactly what can be derived by repeated applications of materials in \mathfrak{L}_G is something itself captured by what’s in \mathfrak{L}_G . So idealization isn’t available to us here (see Woods (forthcoming-a) for detailed discussion of these cases for whole-theory comparisons.)

0.5 AVOIDING THE DEGREES OF FREEDOM PROBLEM

Moving to the more detailed account of prediction and coherence doesn’t alleviate our problem—the problem is real. Wright (1986) argues on this basis that at least *some* of logic needs to be treated as a priori and fixed for the purpose of carrying out reflective equilibrium. As pointed out by Ahmed (2000), among others, this seems too strong. Even for the Quinoid view that every belief, including those about prediction and coherence, is revisable, it’s enough that there be *some* situation in which it’s reasonable to revise it. As pointed out by Elstein (2007, fn. 4), we would also really like the ability to *explain* when it is and when it is not appropriate to revise in addition to “in principle” revisability for all our logical beliefs. This combination would be a potent solution to Wright’s problem.

We might worry, though, about whether we can manage something so strong as explanation. We can put this issue to the side—though I’m sympathetic—as what seems absolutely unavoidable is some additional, epistemically significant, *guidance* for our revising. If we can’t find this, then Wright’s problem is simply unavoidable. But providing guidance in our revisions seems to again involve mustering additional resources which themselves are revisable. Care needs to be taken to ensure we’re not just moving the bump in the rug. What kind of guidance is necessary and what would it involve?

0.5.1 GUIDANCE

The requisite guidance can’t simply be something as epistemically neutral as a disposition to revise this way as opposed to that, not unless we’re willing to wholeheartedly endorse a purely coherentist notion of epistemic justification. Doing so would be pyrrhic since our problem arises exactly from the existence of ways of restoring “coherence” which deny blatant implications of our general logical beliefs in conflict with our logical intuitions.

Coherentist pictures look far more stable when the background logic used to define a notion of coherence taken as fixed. When even what counts as coherence is up for grabs, they look a bit besides the point. Leaning on what we’re disposed to do in the face of significant recalcitrance doesn’t help much; why

should our temptations be epistemically significant?²⁸

We need guidance which tracks the intuitive reasonableness of certain ways of revising our logical beliefs. We need a way of guiding our revisions which itself helps to convey justification on our revising. But what would such a notion look like? It's difficult to see without compromising on our assumed anti-exceptionalist scruples. There are two natural suggestions which I'll now briefly explore with the implicit (well, now explicit) suggestion that all other alternatives collapse into one or the other.

First, the anti-exceptionalist could expand on the theoretical virtues used in deciding how to revise a particular instance of recalcitrance.²⁹ Revising bit of \mathcal{L}_G or $\mathcal{L}_{p/np}$ is one thing, but revising \Rightarrow or \perp seems significantly ad hoc. If *minimal ad hoc ness* is a theoretical virtue, these revisions will be disfavored.³⁰ Care needs to be taken with this method to preserve both the anti-exceptionalist standpoint and the advantage reflective equilibrium has over whole-theory comparisons.

Second, the anti-exceptionalist could invoke the obviousness of various implications (along the lines of Harman (1986) and Field (2009)) in order to guide our revisions. Or, similarly, they could define a notion of the "depth" of certain revisions and then use this to decide in a principled way which of the various members of \mathcal{L}_G and $\mathcal{L}_{p/np}$ to revise. As with our first patch, care needs to be taken with this method to preserve both the anti-exceptionalist standpoint and the epistemic significance of the favored materials guiding revisions.

Each approach invokes additional materials in order to guide our revisions in a way which avoids the degree of freedom problem. The question then arises whether they retain the distinct advantages of the piecemeal approach without either collapsing into an implausible coherentism or abandoning anti-exceptionalism. I claim, and will now argue, that the plausible ways of developing this kind of guidance do exactly this. It's thus dubious that reflective equilibrium is significantly better than whole theory comparison as an anti-exceptionalist methodology. We'll take the patches in reverse order.

²⁸It's important to be clear about what, exactly, this point amounts to. I'm not denying that the contents of, say, \mathcal{L}_G could explicate a reasonable notion of coherence at a particular point in time. Many reasonable conventionalist and fellow traveler lines will accept something like this. Rather, I'm denying that moving from \mathcal{L}_G to \mathcal{L}'_G , guided merely by what we're disposed to do, is seen as an epistemically responsible way of proceeding. I hope to revisit these interesting issues elsewhere.

²⁹Though finding the right set of theoretical virtues for logic is itself a rather tricky matter. See Russell (forthcoming) for useful examples involving theoretical strength.

³⁰I'll focus on ad hocery below but the points generalize.

0.5.2 THEORETICAL VIRTUES TO THE RESCUE?

Again, reflective equilibrium can be construed as a particular form of whole theory comparison making use of only one theoretical virtue—coherence between \mathfrak{L}_G and $\mathfrak{L}_{p/np}$ —and a limited field of logical alternatives—those easily reachable by piecemeal revisions of our logical beliefs. Given this, adding guidance by adding additional theoretical virtues is attractive. But which theoretical virtues? This is important since theoretical virtues which depend on an implication relation don't provide any advantageous guidance since we could then modify whether or not our current stage of revision has that virtue by modifying the implication relation encoded in \mathfrak{L}_G . Additionally, and because of this, the oscillation problem raised above reemerges since these virtues might themselves trade off against one another in ways which are essentially dependent on our currently accepted implication relation.

Which theoretical virtues have epistemic significance, yet aren't obviously beholden to an implication relationship? The natural ones to lean on are the *aesthetic* theoretical virtues like elegance and minimizing ad hoc-ery.³¹ And, sure enough, revising \Rightarrow or \perp nearly always seems ad hoc. However, even ad hoc-ness potentially needs theorizing in terms of implication. Here's Thagard's gloss:

We cannot condemn a theory for introducing a hypothesis to explain a particular fact, since all theorists employ such hypotheses. The hypotheses can be reprehended only if ongoing investigation fails either to uncover new facts that they help to explain or to find more direct evidence for them. (Thagard 78: 87)

Since we can't always find this in advance, presumably our assessment of whether a move is ad hoc reflects our confidence that the hypothesis won't *directly* explain anything else or be explained by something more fundamental. That is, we're going to need further beliefs about what will or will not be explained by a hypothesis. These beliefs in turn explain whether we see a particular hypothesis as ad hoc.

But explanation and, of course, evidential relations themselves are conditioned by a background implication relationship. This yields our familiar pair of problems. First, we can modify our evaluation of ad hocery by means of modifying our beliefs about \Rightarrow and \perp , thereby changing what can explain what. This seems to undermine the role of ad hocery as a piece of epistemically significant guidance.

Second, and because of this, we can end up in revision cycles when the background logic changes. Criteria like the coherence of our two sets of beliefs and avoidance of ad hocery can trade off one for the other. Add enough

³¹See Lipton (2004: 66) for a gloss on these virtues as aesthetic.

subsidiary members to \mathfrak{L}_G and our previously seeming ad hoc change will look well-evidenced, but massively incoherent; revise \Rightarrow or \perp enough to fix this problem, our theory will look coherent but ad hoc. It's easy to see how to construct a cycle out of these materials.

Similar problems emerge for other choices of theoretical virtues. Unless a virtue is independent of our background logic, the resulting picture become open to the sort of cycles that undermine its effectiveness as a methodology. And if they are dependent, then it seems difficult to treat them as epistemically probative. After all, we've already got a set of basic intuitions about validity and invalidity ($\mathfrak{L}_{p/np}$) playing a role in our decision process. Adding to these something like a psychological or perceptually-based theoretical virtue seems unlikely to make the situation better in any significant way.

Suppose we added something like beauty or elegance to our evaluation method. It's hard to see the resulting method would avoid the degrees of freedom problem. After all, though it seems a bit ad hoc to deny instances of \perp or \Rightarrow , it does have a certain iconoclastic charm to it. And it's entirely unclear that our intuitions about the inelegance or ugliness of a revision would persist after performing the revision.

Even if there is some way to make good on some additional theoretical virtues, it still seems as if we can iterate the degrees of freedom problem here. The guidance we have about how to revise is *itself* open to revision; we can simply, in the face of recalcitrance, revise our judgment about the simplicity, strength, or elegance of a particular proposed revision.

Why is this? Just like with our beliefs about \Rightarrow and \perp , we're going to have to lean back on some kind of implicit theory of beauty or elegance to avoid cases where our judgment is simply perverse. But, once we've done this, we can simply revise this implicit theory exactly as we could revise our implicit theory of \Rightarrow and \perp .³² So without additional materials, the degrees of freedom problem in fact grows worse here; more virtues, more degrees of freedom. Constraint is necessary.

If this approach is to be defended, the burden seems clearly on its proponent to give a set of theoretical virtues which manages to actually avoid the degrees of freedom problem, but which doesn't end up reducing reflective equilibrium to cyclic-prone whole theory comparisons. If this is the result of adding theoretical virtues, a case can be made that the proponent of reflective equilibrium needs to adopt constraints in order to block the possibility of cycles and constrain the theoretical options open to the logical revisionary. They then owe us a story about such constraints which (a) stays true to our anti-exceptionalist starting

³²You might think that this implicit theory is independent of our logical beliefs—maybe so. But what it *implies* is not independent of our logical beliefs.

point and (b) isn't open to the selfsame problems. We now turn to our other general option for patching up reflective equilibrium.

0.5.3 IMMEDIATELY OBVIOUS IMPLICATION TO THE RESCUE?

The solution initially suggested by Wright (and toyed with by Shapiro), involves treating a portion of the logic *used* in logical revision as *a priori* and unrevisable (at least at the present moment). This particular solution abandons anti-exceptionalism, but there is a related alternative that seems *prima facie* more promising. We could use of the notion of an *obvious* or *immediately compelling* implication. This isn't anti-anti-exceptionalist—after all, some non-logical transitions are themselves obvious or immediately compelling. Just consider "today is Tuesday, so tomorrow is Wednesday."

The idea would be, then, that many instances of \Rightarrow and \perp are usually immediately compelling in a way that many other members of \mathfrak{L}_G are not. If we require our successive revisions minimize revisions of obvious or immediately compelling logical beliefs, then we have a principled reason to avoid revising \Rightarrow or \perp , solving the degrees of freedom problem without abandoning our anti-exceptionalist starting point.

The trouble is that it's unclear why transitions that seem obvious should be somehow epistemically privileged over those which do not. After all, many of our implicit balks at valid implications are rather deep. Likewise, much of the appeal of fallacious reasoning resists even expertise about the underlying issues, as the Wason selection test (Wason 1968) and Monty Hall problem seem to show. Why then take a property of some transitions which regularly misfires to be of epistemic significance? If we accepted a flatly coherentist picture of epistemic justification, then perhaps such psychological guidance is rationally permissible, but it's exactly this sort of contingency in acquiring a coherent set of beliefs that puts people off coherentism in the first place.

Moreover, few of the tensions motivating our actual logical theorizing are cases of obvious or immediate implications conflicting with other parts of our logical beliefs—these tend to be screened out relatively early. The more typical situations are (a) distinct non-obvious parts of our logical beliefs conflicting with each other and (b) conflict between beliefs which are more or less equally obvious. But in such cases we really don't have sufficient guidance to resolve the problem. So there's a large conjecture for this approach—that \perp and \Rightarrow really are typically more obvious than other revisions—which we need to make good on. I'm rather pessimistic.

Putting such problems to the side, this kind of approach privileges $\mathfrak{L}_{p/np}$ over \mathfrak{L}_G in ways that unbalance the point and anti-exceptionalist character of reflective equilibrium. After all, if we refrain from revising away our disposition to treat various implications as obvious or immediate, then we treat our

dispositional logical intuitions as privileged over our background logical theory. Given the intuitive parity between average members of $\mathfrak{L}_{p/np}$ and those dealing with immediacy, this means that we'll be guided to revise \mathfrak{L}_G at the expense of $\mathfrak{L}_{p/np}$, treating our intuitions about logical theory as having more weight than settled bodies of logical theorizing. This seems problematic; why should our dispositional beliefs about logic—especially in light of the well-known problems for them—be treated as epistemically privileged?

However, if we treat our dispositions to recognize various members of \mathfrak{L}_G as obvious or immediate as *as* revisable as any other belief about logic, then we lose our solution to the degrees of freedom problem; We can improve coherence by moving our dispositional belief that an obvious implication is obvious from \mathfrak{L}_p to \mathfrak{L}_{np} as well as revising the implication itself. One might claim that making multiple revisions should be resisted, but as pointed out above, revisions of our logical beliefs are nearly always going to require making adjustments elsewhere in \mathfrak{L}_G and $\mathfrak{L}_{p/np}$. So this way out looks problematic, resulting in either an epistemically unhealthy dependence on contingent psychological dispositions or resulting in an epicycle of the degrees of freedom problem.

0.5.4 GENERALIZATION AND UPSHOT

As we've just seen, these patches are flawed. They each fail in at least one of four ways, depending on the particulars of the view. First there are patches which add guiding materials which aren't revisable:

- If the requisite guiding materials are of the right type to be revised—say, beliefs about elegance or simplicity—then these patches simply exempt them from revision. But this amounts to a form of exceptionalism, albeit not quite exceptionalism about logic. It seems out of the spirit of the anti-exceptionalist program to allow guiding materials to be immune to revision but allow logic to be revised. It also seems to privilege the less epistemically fundamental over the more epistemically fundamental, losing a bit of the justificatory credentials of the method as well.
- Alternatively, we might add guiding materials which are by their nature, say, resistant to revision. These might include our psychological inclination to revise this way rather than that. But this is problematic since it's unclear why our psychological dispositions are probative, especially given that we've already accounted for them once by their inclusion in \mathfrak{L}_p .

Second there are patches where we don't exempt the guidance from revision:

- But if we can revise our judgments about, say, the elegance or simplicity of a particular revision, the degrees of freedom problem remains. We can just simultaneously revise \Rightarrow and our judgments about the inelegance of so revising in any case of recalcitrance. In principle, adding revisable materials makes the degrees of freedom problem worse, not better. More materials, more degrees of freedom.

And for many, if not most, of these approaches:

- Adding in a second theoretical virtue, such as elegance, makes oscillation more likely as these virtues can trade off against one another; since our evaluation itself might change as we revise, we'll risk being in a position where the virtues of our previous logic and the vices of our current one become apparent once we've revised. Again, adding in more theoretical virtues tends to undermine the putative advantage of piecemeal methods like reflective equilibrium in dealing with problems like oscillation.

So it looks like adding enough guidance to avoid the degrees of freedom problem can't be done without either obviating the putative advantages of the piecemeal method or relinquishing its justificatory credentials. At best, the burden is on the defender of the piecemeal approach to find an anti-exceptionalist way of guiding our revisions which is both epistemically respectable and not open to iterations of the degrees of freedom problem.

There doesn't seem to be a reasonable way to salvage reflective equilibrium as a distinctive methodology for logical revision. In short, the advantages the piecemeal approach seemed to have over whole-theory comparisons comes at simply too high a cost. But the piecemeal approach does track methods we actually use in carrying our revisions to our logical beliefs sometimes, so it's worth briefly closing with what the most reasonable version of the piecemeal approach looks like.

0.6 LEARNING FROM FAILURE

I conjecture that the best option for the anti-exceptionalist inclined towards piecemeal revision involves adding two elements. First, they should make use of additional theoretical virtues, such as the ad hoc ness of various resolutions of tension, to increase the *guidance* of their choice of coherence-increasing adjustments. Unfortunately, as pointed about above, this only helps if these theoretical virtues are themselves such that revisions of them are trayfe. And just stipulating this again seems to do disservice to the anti-exceptionalist position. So we need to find a way of blocking such revisions which is in the spirit of anti-exceptionalism.

How do we block these revisions? One way of doing so, my favored, would involve adopting a constraint I've defended elsewhere. This constraint, LOGICAL PARTISANS, says that we only ought to revise our logical beliefs when both my current viewpoint and the revised viewpoint agree that the revision wasn't any worse than not revising. It's original motivation was to minimize the danger of revision cycles, but it has a payoff in this context.³³

³³LOGICAL PARTISANS doesn't entirely avoid this danger for reflective equilibrium since it's possible that a sequence of adjustments, each locally acceptable, would result in a cycle. LOGICAL PARTISANS can be generalized for sequences of revisions, of course, but perhaps it would be better to just live with a little epistemic risk. I hope to discuss this in further work.

Suppose that there’s an opportunity to revise, albeit ad hoc ly, \Rightarrow in order to resolve some tension. Let T_1 be the current state of our logical beliefs (i.e. \mathfrak{L}_G and $\mathfrak{L}_{p/np}$) and T_2 the state that results from revising \Rightarrow and making the other necessary adjustments. T_2 , though it resolves the tension with T_1 , nevertheless incurs a cost in the measure of its ad-hoc-ness. So T_2 will typically be less preferable than other resolution. This initially promising bit of guidance requires nothing from LOGICAL PARTISANS.

LOGICAL PARTISANS, though, blocks the lurking alternative of revising both \Rightarrow and our judgments about how ad hoc revising \Rightarrow is. This revision can be done either by adding to \mathfrak{L}_G materials that predict the change, thereby reducing the ad hoc character of the change or by simply revising our judgment about how ad hoc T_2 is. However, even though T_2 will now judge T_2 preferable to T_1 ; by the lights of T_1 , T_2 is still massively ad hoc, which is a strongly bad-making feature of a logical view. So, so long as ad hoc-ness is weighted sufficiently heavily, T_2 is worse by both the lights of T_1 . LOGICAL PARTISANS thus blocks this potential move.³⁴ We have the guidance we need to avoid the degrees of freedom problem in this case.

Similar stories can be told about revising \perp and the theoretical virtues themselves. The story just told captures the necessary amount of guidance demanded by the degrees of freedom problem without abandoning the piecemeal project entirely. Our current story generalizes whole-theory comparisons to a wider class of alternatives where we don’t have a settled body of information about the virtues and vices of the logics being compared. It also gives a diachronic, anti-exceptionalist, justificatory story about how we might eventually end up at a settled position about logic.³⁵

Summing up, the combination of LOGICAL PARTISANS and theoretical virtues other than coherence will plausibly block the most troubling cases of the degrees of freedom problem. It does so without eliminating the main virtue of the piecemeal method—the freedom to revise our logical beliefs in different, equally coherence-increasing ways. It favors revision of the non-metatheoretic bits of \mathfrak{L}_G or $\mathfrak{L}_{p/np}$ —i.e. not the bits concerning coherence, prediction, and the-

³⁴The weighting itself helps to resolve the problem mentioned above about which logic we use in predicting coherence; LOGICAL PARTISANS requires using both T_1 and T_2 in evaluating moving to T_2 from T_1 . If the theoretical virtues are heavily weighted enough, then grounds for revising our judgments about them (as well as about prediction and coherence) will be sufficiently demanding that the recovery of some tension when we evaluate T_2 by the lights of T_1 won’t automatically be problematic. One more detailed option is to think of us as generating a field of options $T_n \dots T_k$ using T_1 , then evaluating whether each pair (T_1, T_j) violates LOGICAL PARTISANS. A story would still need to be told about how much “closure” we build into \Rightarrow when proceeding with reflective equilibrium, but that’s not an issue I can enter into in this essay.

³⁵Providing a justificatory adjunct to stories like Russell’s (2014) about how to come to have a particular set of logical beliefs.

oretical virtues—since such revisions will typically not result in dramatic shift of the theoretical goodness of our logical theories. It remains staunchly anti-exceptionalist, as no principle, intuition, belief about validity, or belief about the theoretical virtues of a logical system is *in principle* blocked; it contains, however, sufficient exceptionalism *about our current standpoint* to largely avoid both the degrees of freedom problem and the oscillation problem. My closing conjecture is that this combination of LOGICAL PARTISANS and multiple theoretical virtues is the best way to address Wright’s problem and, thereby, develop the piecemeal method for logical revision.

References

- Ahmed, A. 2000. Hale on Some Arguments for the Necessity of Necessity. *Mind* 109(433): 81–91.
- Bueno, O. and M. Colyvan. 2004. Logical Non-Apriorism and the Law of Non-Contradiction. *The Law of Non-Contradiction: New Philosophical Essays*: 156–175.
- Church, A. 1964. The history of the question of existential import of categorical propositions. In *Logic, Methodology, and Philosophy of Science (Proceedings of the 1964 International Congress)*, ed. Y. Bar-Hillel. North-Holland.
- Elstein, D. Y. 2007. A New Revisability Paradox. *Pacific Philosophical Quarterly* 88(3): 308–318.
- Field, H. 2009. What is the Normative Role of Logic? *Aristotelian Society Supplementary Volume* 83(1): 251–268.
- Goodman, N. 1983. *Fact, Fiction, and Forecast*. Harvard University Press.
- Hanson, W. H. 1989. Two Kinds of Deviance. *History and Philosophy of Logic* 10(1): 15–28.
- Harman, G. 1986. *Change in View*. MIT Press.
- Hjortland, O. T. 2017. Anti-exceptionalism about logic. *Philosophical Studies* 174(3): 631–658.
- Lipton, P. 2004. *Inference to the Best Explanation*. Routledge.
- Mares, E. 2014. Belief Revision, Probabilism, and Logic Choice. *Review of Symbolic Logic* 7(4): 647–670.
- McGee, V. 1985. A Counterexample to Modus Ponens. *Journal of Philosophy* 82(9): 462–471.
- McPherson, T. 2015. The Methodological Irrelevance of Reflective Equilibrium. In *The Palgrave Handbook of Philosophical Methods*, 652–674. Springer.
- Meadows, T. and Z. Weber. 2016. Computation in Non-Classical Foundations? *Philosophers’ Imprint* 16(13).

- Priest, G. 2006. *Doubt Truth to be a Liar*. Clarendon Press Oxford.
- Priest, G. 2016. Logical Disputes and the *a priori*. *Logique et Analyse* 59(236): 347–366.
- Quine, W. 1970. *Philosophy of Logic*. Harvard University Press.
- Rawls, J. 1971. *A Theory of Justice*. Harvard University Press.
- Resnik, M. D. 1997. *Mathematics as a Science of Patterns*. Clarendon Press Oxford.
- Russell, G. 2014. Metaphysical Analyticity and the Epistemology of Logic. *Philosophical Studies* 171(1): 161–175.
- Russell, G. 2015. The Justification of the Basic Laws of Logic. *Journal of Philosophical Logic* 44(6): 793–803.
- Russell, G. forthcoming. Deviance and vice: Strength as a theoretical virtue in the epistemology of logic. *Philosophy and Phenomenological Research*.
- Scanlon, T. M. 2014. *Being Realistic About Reasons*. Oxford University Press.
- Shapiro, S. 2000. The Status of Logic. *New Essays on the A Priori*: 333–366.
- Shapiro, S. 2014. *Varieties of Logic*. Oxford University Press.
- Wason, P. C. 1968. Reasoning about a rule. *Quarterly Journal of Experimental Psychology* 20(3): 273–281.
- Williamson, T. 2014. Logic, Metalogic and Neutrality. *Erkenntnis* 79(2): 211–231.
- Woods, J. 2018. Mathematics, Morality, and Self-Effacement. *Noûs* 52(1): 47–68.
- Woods, J. forthcoming-a. Logical Partisanship. *Philosophical Studies*.
- Wright, C. 1986. Inventing Logical Necessity. *Language, Mind, and Logic*: 187–209.