

The Epistemic Normativity of Logic

by

Kieran Britt

B.A., University of Colorado, Boulder, 2020

A thesis submitted to the Faculty of the Graduate School of the University of Colorado in partial fulfillment of the requirement for the degree of Master of Arts

Department of Philosophy

2021

Committee Members:

Graeme Forbes

Graham Oddie

Matthias Steup

Abstract

Britt, Kieran (Master of Arts, Philosophy)

The Epistemic Normativity of Logic

Thesis directed by Professor Graeme Forbes

This master's thesis investigates the (epistemic) normativity of logic for belief and for reasoning. It begins by outlining what a logic is, and then provides a general framework for providing theories of epistemic normativity. It then investigates the normativity of logic for belief, and draws a set of plausible bridge-principles which link the facts of logic to normative facts about belief from recent work in the literature. Then, the normativity of logic for reasoning is discussed: an overview of theories of reasoning is given, and a sketch of an original theory of reasoning is given. As a result of this theory, it is argued that logic provides constitutive norms for deductive reasoning. Finally, it is investigated how the primary norms of logic for belief relate to evidentialism. It is argued that the norms of logic for belief are strictly incompatible with evidentialism, and that we should nevertheless accept the normativity of logic, opting for a weaker form of evidentialism.

Acknowledgements

This thesis represents the culmination of my time at the University of Colorado Boulder, first in the BA program in philosophy and linguistics, and then later in the accelerated master's program in philosophy. I could never have imagined writing a thesis like this when I first declared a major in philosophy four years ago. It's hard to express just how important of a role my various mentors played in my development over the time I've spent here at CU. I was very fortunate to have my three most important mentors, Graeme Forbes, Graham Oddie, and Matthias Steup, as my committee on this project. So I'd like to take a proverbial moment to thank each of them individually for all their mentorship over the years.

I'd like to thank Graeme for his mentorship throughout my time here, particularly in logic and philosophy of language. Most of what I know in those disciplines I either learned from Graeme directly or absorbed by observing him. I still struggle to emulate the attention to technical detail that Graeme exhibits in his work, and whenever I do any kind of technical work, I do so with the standard in mind of doing work that would (hopefully) get past him unscathed. This has made me a much better philosopher than I otherwise would be. I can only hope to one day have the depth of technical knowledge that Graeme possesses.

I'd like to thank Graham for his mentorship in philosophy generally. Just about everything I know about probability, formal value theory, and metaphysics I learned from Graham, as well as a fair amount of logic, and more besides. More importantly, Graham has taught me an immense amount about general philosophical methodology, and through his own example has given me a model of how to do philosophy that I aspire to emulate.

I'd like to thank Matthias for his mentorship, particularly in epistemology. Matthias's instruction in epistemology is the basis of a great deal of my current philosophical expertise, and

his instruction has greatly influenced both my views and how I tackle epistemological issues (all for the better, as far as I can tell). Matthias also was the one who taught me how to dive deep into the philosophical literature in an area; a skill which has proved invaluable. And he had the greatest hand in shaping me into a competent philosophical writer, through his meticulous and helpful comments on many different papers over the past few years. This is of course not to say that he is in any way responsible for any lack of clarity in this thesis.

In addition to these three mentors, I also owe thanks to:

- Jason Potter, for his mentorship and encouragement early in my philosophical career, which started me down the path of pursuing philosophy seriously.
- Julia Staffel, for teaching her seminar on rationality and reasoning in Fall of 2020. That seminar directed me to most of the literature which I used in this thesis, and my final paper for that seminar ultimately became the core of this thesis.
- My parents, for all their love and support.
- My wonderful girlfriend Angelique, who stood beside me throughout this terrible pandemic, and throughout the process of writing this thesis, even when I was too busy to give her the time and attention she deserved. I don't know how I would have survived this past year without her.

Contents

Introduction.....	1
Part 1: What is Logic?.....	4
A Realist Conception of Logic.....	4
Characterizing Validity and Logical Consequence.....	8
The Deductive Rules of Logic	11
Conclusion to Part 1	13
Part 2: Epistemic Normativity	14
A Sketch of an Epistemic Framework	18
Evidentialism	22
Part 3: The Normativity of Logic for Belief	24
Three Kinds of Norms	26
Bridge Principles.....	29
Desiderata for Bridge Principles	33
The Correct Bridge Principles.....	43
Logical Truth, Logical Falsehood, and the Norms of Logic.....	46
Clutter Avoidance and Pure Epistemic Norms	50
Conclusion to Part 3.....	54
Part 4: The Normativity of Logic for Reasoning	56
What is Reasoning?.....	56
How Does Reasoning Fit into a Theory of the Epistemic?.....	62
A Sketch of a Theory of Reasoning	68
The Advantages of This View	77
The Norms of Reasoning and Epistemic Norms Revisited.....	81
Conclusion to Part 4.....	86
Part 5: Integrating the Normativity of Logic into a Theory of Epistemic Normativity	88
Recap: The Norms of Logic.....	88
What Kind of Norm is Evidentialism?.....	90
Evidentialism and Steinberger’s Directive	94
Is Steinberger’s Directive Reducible to Evidentialism?	98
The Case For NR-	99
A Brief Aside: The Hyperintensionality of Evidence	105

Returning to the Case for NR-	107
The Case For NR+	110
Evidentialism and The Normativity of Logic for Reasoning.....	114
Final Conclusion	116
Bibliography	118

Tables

Table 1: Summaries of Bridge-Principle Desiderata	37
Table 2: Bridge Principle Desiderata and the Norms to which they Apply	38

Introduction

This master's thesis is on the normativity of logic. Specifically, it is on the *epistemic* norms that logic generates for belief and for (theoretical) reasoning. While it seems to me quite plausible that, if logic generates epistemic norms, it likewise generates *practical* norms on practical reasoning and perhaps the formation of intentions, my concern in this thesis is purely with the epistemic.

This thesis is divided into five parts. In Part 1, I briefly discuss the nature of logic, and articulate my view about what formal logic is the study *of*. In brief, I adopt the standpoint of what would plausibly be considered to be a kind of *realism* about logical facts: there are genuine logical facts that are in some sense “out there” about the logical structure of natural language and/or the world. On this view, the various formal logics are best construed as *theories* about logic. These theories will be accurate/inaccurate depending on how closely they describe the actual logical structure of the world or at least natural language. Thus, to say that logic is epistemically normative is to say that the facts of the true logical theory (whichever one that is) generate norms on belief or on reasoning.

In Part 2, I provide a primer on epistemic normativity and what an overall theory of the epistemic would look like. Since I will be ultimately arguing that logic generates epistemic norms on reasoning and belief, Part 2 is useful in that it describes the role that norms on reasoning and belief play in such a theory.

In Part 3, I give an overview of the literature on the normativity of logic, drawing on classic work by Gilbert Harman and John MacFarlane, and drawing very heavily on recent work by Florian Steinberger. Using Steinberger's work as my guide, I argue that logic generates

certain epistemic norms on belief. I then tweak the most important of these norms to be compatible with certain “purist” epistemic views; whether logic is normative for belief should not hinge on the falsity of such views.

In Part 4, I discuss the norms of logic on *reasoning*. I first overview the literature on reasoning, and settle on a view about what reasoning is, drawing from the work of Paul Boghossian and others. I then investigate the relationship between the rules of reasoning and the rules of belief revision. I pose some problems for the view, held by Gilbert Harman and many others, that the rules of reasoning just are the rules of belief revision. I then sketch a novel account of reasoning, drawing insight from the work of Pavel Tichý, and argue that the rules of reasoning are rules of belief formation with regard to a special class of propositions which I call *support statements*. I show how my view alleviates the problems posed earlier, and discuss how my view would fit into an overall theory of the epistemic. The upshot for the normativity of logic is that, on my view, the rules of deductive logic set constitutive norms on deductive reasoning, and determine which rules are acceptable: namely those that correspond to the rules of deductive logic that one is justified in employing. My view thus vindicates a very robust form of the normativity of logic for reasoning, and repudiates the view infamously held by Gilbert Harman that the rules of logic are of no special relevance to reasoning.

In Part 5, I discuss how the norms of logic defended in Parts 3 and 4 fit into an overall epistemological theory of the sort discussed in Part 2. In particular, I focus on whether the main norm I defend in Part 3, which I call ‘Steinberger’s Directive’, is compatible with Evidentialism. After surveying the options for how they may relate, I argue that the norms of logic on belief are *not* reducible to evidential norms on belief. I then discuss what this means for our overall epistemic theories. The most important upshot is that, if Steinberger’s Directive is true,

Evidentialism is – strictly speaking – false. That said, the overall epistemic picture that we get with Steinberger’s Directive included is very close to the picture defended by Evidentialists. And even if we choose to reject Steinberger’s Directive in favor of Evidentialism, my account of the logical norms on *reasoning* is compatible with Evidentialism.

Taken as a whole, this thesis is thus an extended defense of the epistemic normativity of logic, and an investigation into how the norms of logic can fit into our overall epistemological theories. If I am correct, then logic is normative for reasoning in a very robust sense. There is also a strong case to be made that logic is normative for belief, though this would mean that Evidentialism is false.

Part 1: What is Logic?

In this section, I discuss the nature of logic, and articulate a kind of minimal realism about logic, according to which, there are genuine descriptive logical facts about natural language, and competing logics are best construed as *theories* of these facts. I then discuss the options for characterizing the core logical notions of validity, derivability, and logical consequence. Since this is a thesis about the normativity of logic, it is worth articulating what I take logic to be, and what I mean when I say that the “facts of logic” are normative for belief or reasoning. The different ways of characterizing validity also plays a subtle role in my account of the normative role of logic for reasoning that I give in Part 4.

A Realist Conception of Logic

Despite the enormous amount of philosophical work that is either on logic itself or that applies logic to philosophical issues, there is no consensus about what logic is. Many authors define logic as the study of reasoning, though as we will see, this is highly contentious.¹ Perhaps the least controversial definition of logic is that logic is, “the study of the difference between *valid* and *invalid* arguments”.²

What is a valid argument? The pre-theoretic idea is this: some arguments are good, and others aren't. “Every man is mortal, Socrates is a man, therefore Socrates is mortal” is good, while “Every man is mortal, Socrates is mortal, therefore Socrates is a man” isn't. The first is good because its conclusion, in some (specifically logical) sense, *follows* from its premises. That is, the conclusion couldn't fail to be true if the premises were true. Similarly, the second is bad

¹ See, e.g., Hedman 2004, p. xiv.

² Forbes 1994, p. 1.

because in it, the conclusion doesn't, in that same sense, *follow* from the premises. What's more, the reason the first argument is good but the second isn't has to do with the "structure" or "form" of the sentences. In a slogan, a valid argument is "necessarily truth preserving, in virtue of its form".³

It will be worth drawing attention to a couple of simple observations. First, the validity of the first argument (and invalidity of the second) is a descriptive fact. Minimally, it's a fact about *natural language*. While we can make precise exactly why the argument is valid (invalid) by translating it into a formal language, the fact of the argument's validity (invalidity) doesn't hinge on any facts about a particular formal language that we could translate the sentences into. Second, the validity or invalidity of each argument plausibly depends, at least in large part, to the meaning of certain parts of the sentences. Just as,

(1) "Every man is mortal, Socrates is a man, therefore Socrates is mortal"

is valid, so is

(2) "Every dog is kind, Fido is a dog, therefore Fido is kind".

But,

(3) "Some man is mortal, Socrates is a man, therefore Socrates is mortal"

is not valid. Why does substituting expressions of the same syntactic type preserve validity in the move from (1) to (2), but not in the move from (1) to (3)? It's because the vocabulary we changed in moving to (2) was in some sense *non-logical*, while the vocabulary we changed in moving to (3) was *logical*. The meanings of the logical vocabulary are what determine the truth or falsity of facts about validity in natural language. Of course, it is no easy task to specify what

³ This is close to the pre-theoretic definition of logical consequence given in Sider 2010, p. 2.

makes certain expressions logical and others not.⁴ But I'll take it for granted that *at least* the standard truth-functional connectives, the quantifiers of first-order logic, and the identity relation are uncontroversially logical.

Logic isn't only concerned with validity. There are a family of prototypically logical notions which are often discussed: *validity*, *derivability*, *entailment/consequence*, etc. For each of these notions, there are competing ways of characterizing the notion (e.g., proof-theoretic or semantic). There are also competing views about what kinds of objects the notions apply to; that is what kinds of objects arguments are composed of. Some contenders are *sentences*, *propositions* (construed as sentence-meanings), and *statements* (the utterances of sentences or propositions). And most importantly, there are competing views about which sets of said objects possess said properties. That is, there are competing views about which arguments are valid, which sentences or propositions follow or are derivable from which others, etc.

What I have just said in the last few paragraphs is an articulation of a fundamentally realist conception of logic. There are genuine descriptive facts about which arguments are valid, and which sentences/propositions follow from others. These are facts about natural language *at least*, though I certainly don't mean to rule out that they also may double as facts about the underlying structure of reality.⁵ On this view, the various formal logics should be construed as *theories* of logic: they are attempts to describe the underlying logical facts of natural language, and as such, may be more or less accurate descriptions of the underlying logical facts. These facts are determined by the meanings of the logical constants in natural language. We can call this view *minimal realism*.

⁴ See MacFarlane 2017 for a thorough overview of the various approaches for specifying which expressions are logical.

⁵ For views of this sort, see Sher 2011 and McSweeney 2019.

An anti-realist, by contrast, will hold that there are no such facts. Perhaps there are facts about whether an argument is valid in this or that formal logic, but there is no meaningful answer to the question of whether the argument is valid *simpliciter*.⁶ Such an anti-realist will likely be rather unenthusiastic about the normativity of logic: if there is simply no fact of the matter about whether a given argument is valid *simpliciter*, then it's hardly worth asking if logic generates any norms about how to reason or what to believe.

If one grants minimal realism, then another line of questioning opens up: the question of logical pluralism. Is there a unique logical theory which is correct, or are there multiple correct theories (which are perhaps correct for different things)? To give the former answer is to accept *logical monism*: the view that exactly one logic correctly describes the logical structure of language.⁷ To give the latter answer is to adopt *logical pluralism*. Logical pluralism comes in various forms, some significantly more radical than others.⁸ While there is extensive room for debate regarding logical pluralism, for my purposes in this thesis, I'm going to *assume* it is false. There is good reason for this: as Florian Steinberger has convincingly argued, there is a tension between logical pluralism and the normativity of logic. Steinberger describes the tension thus:

If logic is normative, it tells us something about how we ought to reason. If, as the pluralist would have it, there are several correct logics, those logics make incompatible recommendations as to how we ought to reason. But then which of these logics should we look to for normative guidance? I argue that inasmuch as pluralism draws its motivation from its ability to defuse logical

⁶ It's worth noting that the realism/anti-realism distinction I'm drawing is distinct from another realism/anti-realism divide one could draw with regard to logic: the question of whether logical facts are "in the mind" or "in the world". The former sort of view constitutes a form of psychologism, while the latter is a kind of metaphysical realism about logic. While I firmly side with the latter sort of view, I need take no stand on this question for my purposes in this thesis.

⁷ For a stalwart defense of logical monism, see Tichý 1988, ch. 14.

⁸ For more on logical pluralism, see Russell 2021 and Cook 2010.

disputes—that is, disputes between advocates of rival logics—it is unable to provide an answer: pluralism collapses into monism with respect to either the strongest or the weakest admissible logic.⁹

Since my goal in this thesis is to defend the normativity of logic from general skeptical concerns, and to investigate the epistemological consequences of logic’s normativity, I will assume that logical pluralism is false.

I will thus adopt a realist, monist conception of logic. My concern in this thesis is whether the logical facts about natural language generate norms on belief or on reasoning, and if so, how these norms fit into an overall epistemological theory. While I will assume that some formal logic or other gives the correct description of the underlying logical facts, I won’t take a stand on which one is correct (though many of the epistemological concerns discussed by the authors I cite make the most sense against the backdrop of classical logic). I will, however, assume that logical relations such as logical consequence are relations between *propositions*, construed as the meanings of sentences.¹⁰

Characterizing Validity and Logical Consequence

Given minimal realism, there are two main contenders for how to specify a theory of logic. The two contenders correspond to two different ways one can specify the meaning of the core logical notions, i.e., the logical constants.

⁹ Steinberger 2019c, p. 1.

¹⁰ I won’t endeavor to give an analysis of what a proposition is in this thesis, though I hold the following things to be true. First, propositions are the meanings of sentences. Second, propositions are the bearers of truth values. Third, propositions are *fine grained* (hyperintensional): logically equivalent sentences are not necessarily synonymous, despite being necessarily co-extensional and co-intensional. I’ll discuss hyperintensionality when it becomes important for this thesis much later on.

One way to specify the meaning of the logical constants is to give what is called a *denotationalist theory* of logicity.¹¹ Broadly speaking, on this approach, the core logical notion is that of *validity*. A denotationalist theory starts by giving model-theoretic specifications of the semantics of the logical constants: i.e. of the truth functional connectives and quantifiers (and the two-place identity predicate). An argument consists of a set of premises (sentences/propositions) and a conclusion, and an argument is valid iff the conclusion must hold a *designated* truth value in every model-theoretic interpretation where all the premises hold a designated truth value. For classical logic, this will mean that an argument is valid iff, in any interpretation where the premises are true, the conclusion is true. But I leave it open that perhaps the true logic is many-valued, and might have other values than just *truth* as designated.

Given a denotationalist specification of the meanings of the logical constants, a specification of a system of proof for the logic is secondary, and adequacy of a given proof system relative to the logic is measured first and foremost by whether the proof system is *sound* with respect to the semantics of the logic.

The second way to specify the meaning of the logical constants is to give what is called an *inferentialist theory* of logicity.¹² At its core, an inferentialist theory specifies logicity by first specifying the proof theory of the logic, i.e., the inference rules which are acceptable for each logical constant. In its strongest form, this sort of view takes *derivability* as its core notion, and a conclusion is *derivable* from a set of premises iff one can derive the conclusion from the premises through a finite series of applications of the allowable inference rules. On this view, the semantics of each logical connective is secondary, and the adequacy of a given account of the

¹¹ Denotationalism, while rarely explicitly defended, is far and away the most popular way of characterizing the meaning of the logical constants, and is adopted in just about every logic textbook.

¹² For a thorough overview of inferentialism, see Murzi and Steinberger 2017.

semantics of the connectives is measured by whether the semantics makes each argument whose conclusion is derivable from its premises valid.

Inferentialist accounts of logicity are typically associated with use-theoretic theories of meaning. If one is skeptical of the notion that linguistic competence consists of knowledge of the meaning of expressions, one might instead think that linguistic competence consists in knowledge of when it is acceptable to *use* this or that expression. Applied to the logical constants, this amounts to knowledge of the assertability conditions of sentences, i.e., of when it is acceptable to infer a sentence/proposition from others. Thus, many philosophers who are skeptical of broadly referential theories of meaning are sympathetic to inferentialist accounts of logic.

But one can be an inferentialist without thereby rejecting referential theories of meaning more broadly. On this sort of view, while the inference rules of each logical constant characterize the meaning of said logical constant (and perhaps are what must be grasped for the constant to be understood), the logical constants nevertheless have denotations.¹³ It is this sort of inferentialism I am most concerned with for the purposes of this thesis.

So a denotationist first specifies the semantics of each logical constant, then uses this to specify validity as designated truth-value preservation in all interpretations. They then can give a proof-theory which is sound and (in the case of first-order logic or less expressive theories) complete with respect to their semantics. By contrast, an inferentialist first specifies the rules of inference which are acceptable for each logical constant, which settles the proof theory of the logic. The inference rules for a logical constant also delimits the range of denotations that said logical constant can have, though it doesn't by itself determine it. However, given certain other

¹³ For one such account, see Peacocke 1987.

semantic assumptions (such as assumptions about the truth values a proposition could bear), the inference rules will determine the semantics of the logical constant, thus setting the model-theoretic semantics of the logic.¹⁴

The Deductive Rules of Logic

As it will become important for my account of the normativity of logic for reasoning in Part 4, it's worth briefly discussing how inferentialists and denotationalists of the sorts I discussed in the last section can settle on a particular deductive system.

For an inferentialist, there is a substantial question about which rules of logic are primitive, since these characterize both the logical consequence relation of the logic, as well as the meanings of the logical constants. We might, following Peacocke, take these to be the rules which are “primitively compelling”; i.e., those that must be grasped for a speaker to grasp the meaning of the constant. But there are various other desiderata that can be given, depending on one's view, which will limit which rules are legitimate. We might take it that each constant should have its own introduction and elimination rules. We might also demand that the introduction and elimination rules for each constant be *harmonious*, in that, in introducing a logical constant, we may not allow that we could infer conclusions from the constant's elimination rules that would not be justified by the premises of the corresponding applications of its introduction rules.¹⁵ This is all to say that an inferentialist will have various grounds for

¹⁴ This idea is discussed briefly in MacFarlane 2017, and in more depth in Peacocke 1987. It originates in the “do-it-yourself” semantics of Hacking 1979.

¹⁵ This requirement avoids, for example, the introduction of Prior's operator Tonk (Prior 1960). See Steinberger 2011 for an extended discussion of the motivations behind harmony, and how it may be formulated.

identifying certain inference rules of logic as primitively acceptable, and will be able to then justify further rules in terms of the primitive ones.

At first glance, it might seem that, for a denotationist, any set of inference rules which is sound and complete with respect to the semantics the denotationist adopts will be perfectly acceptable. But even a denotationist can give grounds (of the same sort as the inferentialist) for privileging certain sets of inference rules, especially if they take themselves to be giving a characterization of the rules of logic *simpliciter* (as the realist suggests they are). For example, we might have a general preference for systems of rules over systems of axioms which accord better with deductive practice. We might prefer deductive systems which give rules for each of the logical constants individually, on the ground that an agent who grasped only a subset of the logical constants we grasp should nevertheless be capable of utilizing the full inferential capabilities of the constants he does grasp. For the same reason, we may also wish our rules to be harmonious. Further, we might want rules which match the semantics of each constant we give: a defender of an arbitrary object semantics for quantifiers will favor the Fitch-style instantiation and generalization rules for quantifiers, while someone who takes quantifiers to represent implicit conjunctions (in the case of the universal quantifier) and disjunctions (for the existential quantifier) over the domain will favor the Gentzen-style introduction and elimination rules.¹⁶

So regardless of whether one provides an inferentialist or denotationist stance, it is plausible that one can give grounds for privileging certain sets of inference rules as characterizing the basic deductive system of the logic. Provided that no new logical constants are introduced, further rules can then be justified by deriving them from the primitive ones.

¹⁶ See Fine 1985 for a defense of the former, and Forbes 1993 for a defense of the latter.

Conclusion to Part 1

In this section, I've discussed the nature of logic, and articulated a realist conception of the study of logic, where rival logics are competing theories about the underlying logical structure of the bearers of truth in natural language (whether these are linguistic items like sentences, or non-linguistic items, like propositions). I've also discussed the two available options (inferentialism and denotationalism) for characterizing the meaning of each logical constant, and hence, for characterizing the relation of logical consequence, as well as validity and derivability of arguments. I then briefly discussed the grounds that the inferentialist and denotationalist can give to specify the deductive rules that they accept as primitive. This hopefully should make clear what I take logic to be when I discuss the normativity of logic.

Part 2: Epistemic Normativity

In this section, I will give a primer on epistemic normativity and what a theory of it should contain. I will then sketch a general picture of epistemic normativity that will inform my analyses later on. Since the focus of this thesis is the (epistemic) normativity of logic, it will be worth delineating what epistemic normativity is, and what logic might be (epistemically) normative *for*.

Epistemology is often translated as the “theory of knowledge”, and it is common to think that *knowledge* is the core epistemological notion. Most mainstream epistemological theories typically define knowledge as true belief that possesses some important normative quality or qualities, such as *justification*, *rationality*, *reasonableness*. For instance, many contemporary epistemologists will define knowledge as *justified true belief* plus some fourth condition that is meant to rule out Gettier and relevant-alternative cases.

No matter the definition of knowledge we give, there are questions we can ask about the various normative notions we might attribute to beliefs. For example, we can ask for an account of *justification*. We can also ask for an account of *rationality*, and then ask whether rationality and justification are identical, and if not, which it is that contributes toward a belief counting as knowledge. Different possible answers to these sorts of questions will generate substantively different epistemological theories.

While justification and rationality (and other epistemically normative notions) can apply to beliefs, belief is not the only doxastic attitude to which they can apply. Traditionally, there are three doxastic attitudes: *belief*, *suspension of judgement*, and *disbelief*. We might also entertain *degrees of belief* or *credences*, though I will focus my attention on the traditional doxastic

attitudes throughout this thesis. On most views, doxastic attitudes are attitudes that are taken toward *propositions*, by which I mean nothing more substantive than the claim that what you believe/disbelieve/suspend judgement on is the semantic content of sentences.

Belief is a fairly primitive notion, and it isn't clear to me that it is possible (or useful) to give a definition of it, though I'll take it as uncontroversial that belief is a truth-directed attitude, and that if one believes a proposition, one accepts that proposition to be true. That said, there are significant controversies about the other two doxastic attitudes. First, there is a controversy regarding what suspension of judgement is: some authors seem to take it as the absence of a doxastic attitude, while other authors emphasize that suspension of judgement is the deliberate withholding of belief/disbelief toward a proposition one has considered.¹⁷ I will follow the latter approach, and will thus distinguish suspending judgement on P from simply having no attitude toward P (which may occur because the agent hasn't ever considered whether P). Second, there is controversy regarding the meaning of 'disbelief'. 'Disbelief' is ambiguous in English. On one interpretation, 'disbelief' can mean a mere lack of belief, in which case, disbelief in P will overlap with suspension of judgement on P. On another interpretation, 'disbelief' amounts to explicit *rejection* of P. While this will depend on the semantics one accepts for negation, given the classical semantics, the latter interpretation of 'disbelief' will equate disbelief in P and belief in $\sim P$.¹⁸ Once again, I will adopt the latter interpretation of 'disbelief', and I won't take a stand on whether disbelief in P is identical to belief in $\sim P$.

¹⁷ See Friedman 2017 for a thorough discussion of suspension of judgement.

¹⁸ A dialetheist, by contrast, will reject this equation, since the dialetheist will hold that it is possible to both believe P and $\sim P$, but will not want to thereby say that you can both believe and disbelieve P. There may also be non-dialetheism related reasons to reject the equation of disbelief in P and belief in $\sim P$: bilateralist approaches to natural deduction can be motivated by a distinction between the speech act of *denying* P (corresponding to the attitude of disbelieving or rejecting P) and the speech act of *accepting* $\sim P$. See Rumfitt 2000 and Humberstone 2000 for discussion of this sort of bilateralism, and see Restall Forthcoming, pp. 10-16 for discussion and criticism of this approach.

While providing a clear definition of it is notoriously challenging, there is plausibly a distinctively *epistemic* sort of normativity that is being evoked when we say that you are justified or rational in holding a belief.¹⁹ This sort of normativity is plausibly distinct from moral or prudential normativity: having justified beliefs (whatever that requires) is plausibly distinct from having beliefs which are morally good or in one's own best interest.²⁰ If we grant that the epistemic is a genuine and *sui generis* domain of normativity, then we can make sense of epistemic versions of the kinds of normative notions that characterize norms in any other domain. In particular, we can make sense of and try to define (epistemic) *reasons*, *permission*, *obligation*, *blameworthiness*, *criticizability*, *correctness*, and *fittingness*, to name just a few. For any such notion, there will be substantive issues about how it relates to justification, rationality, or any other candidate property that might link belief and knowledge. With these notions in tow, we can construct distinctively epistemic norms on belief. These norms will make-up the core of a theory of epistemic normativity, and by extension, the core of an overall theory of the epistemic.

The norms on belief I've alluded to thus far have all been *synchronic norms*: we might define what it is for an agent to be, say, epistemically obligated to believe a proposition P *at a given time*. But – somewhat controversially – there might also be *diachronic epistemic norms*. For example, there might be permissible and impermissible ways to *revise* or *formulate* beliefs, and norms on these activities might be distinct from the norms on belief itself. Arguably, any epistemic theory will need to admit of some diachronic norms: it would be hard to make sense of there being, say, a requirement to believe what one's evidence supports at a given time, without

¹⁹ See Cohen 2016 for a sustained argument against the possibility of defining what the epistemic is.

²⁰ Though there are authors who contest this, e.g. Rinard 2019.

there also being a requirement to *formulate* said belief by that time.²¹ But these sorts of diachronic norms will be parasitic on our synchronic norms on belief. It is a more controversial issue whether there could further be diachronic norms on belief formation that do not merely fall out of our synchronic norms on belief. For example, a Bayesian may hold that you are required to update your beliefs by conditionalization on evidence you receive. On this sort of view, the aforementioned diachronic norm is central to the epistemic theory, and is not a mere byproduct of a synchronic norm to have probabilistically coherent credences.

If logic possesses any epistemic normativity, then facts of logic (such as facts about what entails what) will give rise to normative facts regarding either synchronic norms on belief, or diachronic norms on belief formation and revision. Insofar as reasoning is a way to formulate beliefs (or is in some way closely connected to belief formation), any norms on reasoning that logic provides will by extension provide norms on belief revision.

Thus far, I have merely given an abstract overview of what epistemic normativity is concerned with. In the next section, I will provide a general sketch of the kind of epistemic theory I will have in mind when asking about whether logic is normative for belief or for reasoning. Strictly speaking, what I will provide is not a particular epistemic theory, but rather a general framework for formulating epistemic theories.

²¹ Such diachronic norms open up ugly questions about how long one should have to formulate beliefs. But the intuitive idea is that, if I come to learn evidence E which supports belief in proposition P at t, then, since forming beliefs is presumably not an instantaneous process, there is some time shortly after t, t', by which I must formulate a belief in P. There will also then be a requirement to *come to believe* P by t'.

A Sketch of an Epistemic Framework

In the kind of epistemic theory that I have in mind, the most fundamental, primitive normative notion is that of an *epistemic reason*. I'll be taking reasons as primitive, so I won't be giving a definition of what it takes to be a reason. But to get a feel for what I have in mind, we might (unofficially) say that epistemic reasons are the kind of consideration that indicates that a proposition is true.²² We can believe/disbelieve/suspend judgement on propositions, and there are epistemic reasons to do so. These reasons *support* the appropriate doxastic attitude toward the proposition. I'll also take standard normative notions like *obligation* (ought) and *permission* (may/is permitted to) and *criticizability* or *blameworthiness* to be primitively understood, though this is for convenience more than anything: I don't mean to preclude the possibility that we can or should define these notions in terms of reasons.

With the notion of an epistemic reason in hand, we can make the following definitions:

Propositional Justification: An agent S has (propositional) justification for a proposition P at time t iff S's epistemic reasons at t support belief in P.

Doxastic Justification: an agent S is (doxastically) justified in a belief in a proposition P iff:

- (1) S has propositional justification for P
- (2) S believes P
- (3) S *properly bases* their belief in P on the source of their propositional justification for P.

²² Notice that if we took this as an official gloss of 'reason', we would then be left with the notion of a *consideration* as primitive. This gloss is also somewhat controversial, as it builds in a sort of veritism into the notion of an epistemic reason. I take no stand on the issue of what characterizes epistemic reasons here.

Propositional justification is a relation between agents, propositions, and the agent's epistemic reasons. Doxastic justification is a relation between agents and *beliefs* in said propositions, and the epistemic reasons the agent has for the beliefs. I will not distinguish between justification and rationality, and I will assume that *knowledge* is doxastically justified true belief which also satisfies some no-accident condition (i.e., a condition which rules out Gettier cases and the like). For my purposes, the nature of this no-accident condition is unimportant.

The notion of *basing* deserves some explanation. There are many possible accounts of what basing is. We might take basing to be a causal or non-causal relation, though causal accounts are more popular. To base a belief on a reason might be construed as having a certain sort of *disposition* to form said belief when one has said reason, or it might be construed as the application of some sort of *rule* to form said belief. The most sophisticated accounts of basing will usually be hybrids between the two. For our purposes however, we need not specify a theory of basing, beyond saying that to base a belief in P on a reason R is for R to be a reason *for which* one believes P. *Reasons for which* are not merely *reasons why* one believes (they are not mere causes), but further are intended to play some sort of justifying role for the belief.²³

The notion of *proper* basing is meant to recognize that not all ways of basing a belief on a (good) reason are justifying. Rather, for one to be (doxastically) justified in a belief in a proposition, one must further come to believe the proposition *in the right way*. To take an example from John Turri, suppose that you believe $P \rightarrow Q$ and P, and by applying a rule of inference, come to believe Q.²⁴ Suppose further that you have good epistemic reasons (whatever those are) to believe P and $P \rightarrow Q$. As you derive Q from P and $P \rightarrow Q$, you plausibly believe Q

²³ For a thorough and general discussion of the sorts of accounts of basing that are viable, see Neta 2019.

²⁴ I modify the example given in Turri 2010, p. 317.

based on P and $P \rightarrow Q$, as well as your reasons for believing P and $P \rightarrow Q$. If the rule of inference you apply is modus ponens, then it seems that your belief is *properly based* on your reasons. But suppose instead that you form your belief in Q by applying the rule of *Modus Profuscus*: from any propositions A , B and C , infer C from A and B . In this case, even though you base your belief on propositions which you have good reason to believe, and which could doxastically justify your belief in C (in this case, Q), your belief in C (Q) is plausibly not doxastically justified, since you have formed it in the wrong way.²⁵

It's worth noting that this example highlights the minimal role that reasoning (whatever that is) plays in an overall epistemic theory. One way of forming beliefs is to base them on inferences. Some kinds of inferences are justifying, while others are not. If there are norms on correct reasoning (generated by the facts of logic or otherwise) these will fit into the sort of epistemic theory I have in mind by providing criteria for *properly basing* one's inferentially-formed beliefs.

I mentioned earlier that I will take it for granted that we can make sense of other normative notions like *permissibility*, *obligation*, *blameworthiness*, etc. Plausibly, there will be links between these notions and either sort of justification. As a first pass, which will have to be revised later, I'll take it for granted that it would be permissible to believe P iff one possesses propositional justification for P , and that a belief in P is permissible iff it is doxastically justified. To my knowledge, this assumption is uncontroversial.

²⁵ For a more thorough discussion of the connection between propositional and doxastic justification, see Silva and Oliveira 2020.

There is significantly more controversy when it comes to the relation between justification and obligation. There is a live debate in epistemology over whether to accept the following principle:

Uniqueness: For any set of epistemic reasons, there is exactly one attitude which is permissible toward each proposition. Hence, that attitude is obligatory.

If Uniqueness is true, then any permissible (justified) attitude is also obligatory. Uniqueness is often motivated by certain forms of Evidentialism: if you think that your epistemic reasons to take an attitude toward a proposition just are your pieces of evidence for or against that proposition, and you think that your total evidence will always either support, go against, or be inconclusive toward a proposition (but never more than one of those), then you will likely accept Uniqueness. For my purposes, I need not take a stand on the debate regarding Uniqueness, though I will in part 5 investigate the relationship between logic and a form of Evidentialism which accepts Uniqueness.

I've said all of this is simply to sketch the core of the kinds of epistemic theory I have in mind. This will become especially important in part 5, when I discuss how the norms of logic I will defend in parts 3 and 4 are supposed to fit in to our overall epistemic theories. The sketch I've given is relatively non-committal: many different views (of both internalist and externalist persuasions) can fit into the mold I've given, by giving specifications of exactly what counts as an epistemic reason.²⁶ A fully fleshed out theory will also try to specify any other synchronic epistemic norms that exist (what one ought to believe, or what one is criticizable for, etc.) should

²⁶ There are some views, however, which are incompatible with the mold. First, Reliabilism arguably doesn't fit, unless one gives a very strained definition of 'reason'. Views which take doxastic justification to be conceptually prior to propositional justification (see Turri 2010 for an example) also won't fit the definitions I've given.

any of these not be reducible to norms of justification.²⁷ Further, one might also permit more robust diachronic norms on belief formation or perhaps even non-mental actions, such as evidence-gathering. For my purposes, I won't assume any such norms exist, and I will attempt to defend the normativity of logic in parts 3 and 4 with respect to the core synchronic norms that an epistemic theory must specify.

Evidentialism

In part 5, I will discuss the relationship between the norms of logic and other epistemic norms.

One of my main concerns in that section will be how the norms of logic relate to *Evidentialism*.

In principle, Evidentialism can be formulated in either internalist or externalist ways, though the form I am most concerned with is an internalist conception of Evidentialism. We can define

Evidentialism as follows, though I will revise these definitions in part 5:

Evidentialism about Reasons: Epistemic reasons for believing propositions are constituted by an agent's *evidence* for/against those propositions.

Given this definition, we can define propositional and doxastic justification as follows:

Propositional Justification: An agent S has (propositional) justification for a proposition P at time t iff S's (*total*) *evidence* at t supports belief in P.

Doxastic Justification: an agent S is (doxastically) justified in a belief in a proposition P iff:

²⁷ For example, one might think that, in addition to norms to believe according to your epistemic reasons, there are also separate sui generis coherence norms about what it is rational to believe. See Worsnip 2018 for such an account.

- (4) S has propositional justification for P
- (5) S believes P
- (6) S *properly bases* their belief in P on their *evidence* for P.

I will be assuming that evidence is propositional: one's evidence for P are the propositions which one believes/knows which support P. I won't give a more precise account of what it takes for a proposition to be evidence for P (or how much evidence it is) here, though I will briefly discuss this issue in Part 5. I also will be assuming Uniqueness as part of my version of Evidentialism. This means that what one ought to believe will correspond to what one is justified in believing, which yields the following definition:

Evidentialism about the Epistemic Ought: Whether S ought to believe P is determined by whether S has, at the time, adequate evidence for P. Assuming S is aware of the question of whether P, then (1) if S has adequate evidence for P, then S ought to believe P, and (2) if S does not have adequate evidence for P, then S ought to not believe P (i.e., suspend on P or disbelieve P).

With this general epistemic framework in mind, it is now time to turn to the question of what epistemic norms (if any) facts about logic generate for reasoning and for belief.

Part 3: The Normativity of Logic for Belief

Most philosophers and logicians take it for granted that there is a very tight link between logic and reasoning.²⁸ Gilbert Harman has famously cast doubt on this claim.²⁹ For Harman, the rules of reasoning are the rules of “reasoned change in view”, i.e., the rules of *belief revision*. But as Harman pointed out, the rules of logic do not straightforwardly provide rules of belief revision.³⁰ The fact that P and $P \rightarrow Q$ jointly entail Q does not mean one ought to believe Q if one comes to believe P and $P \rightarrow Q$. In some cases, the right thing to do may be to instead reject P or $P \rightarrow Q$. So whatever way logic constrains belief and belief revision, it is not quite as straightforward as an identity between logic and reasoning.

Beyond pointing out the gap between the rules of logic and the rules of reasoning, Harman notoriously argued for a much stronger thesis: that, “there is no clearly significant way in which logic is specially relevant to reasoning.”³¹ On Harman’s view, logic and reasoning are not even importantly related, at least not in any sense deeper than how facts about physics, biology, or what you ate for lunch are related to reasoning.³² Through his arguments, Harman has posed a skeptical challenge to the defenders of the normativity of logic. Anyone who wishes to defend the normativity of logic from Harman’s challenge must attempt to provide a set of plausible norms between facts about logic and facts about reasoning and belief.

²⁸ See just about any introductory logic textbook, for example, Hedman 2004, p. xiv, which *defines* logic as “the study of the principles of reasoning” (ibid) and Hurley 2015, p. xiv, which claims that, “the chief purpose of logic is to develop good reasoning skills” (ibid).

²⁹ Harman 1986.

³⁰ Harman 1986, p. 11-12.

³¹ Harman 1986, p. 20.

³² See Harman 1986, p. 11-20.

Since Harman's influential work, many authors have offered responses to Harman's challenge, with some success.³³ The most developed of such defenses come from two recent papers by Florian Steinberger, where he convincingly argues that there are three distinct normative roles that logic could play, and that there are plausible bridge principles which tie facts about logical consequence to facts about the normative status of beliefs.³⁴

Perhaps surprisingly, Steinberger and other authors who have sought to answer Harman's challenge have typically advocated for principles that link facts about logical entailment to facts about the normative status of *beliefs*. If the rules of reasoning are the rules of belief revision, as Harman thinks, then this perhaps fully settles the question of logic's normativity for reasoning: norms about which beliefs one may hold will naturally place constraints on how beliefs may be revised. But there are other views on what reasoning is, and there are substantive questions about how reasoning relates to belief, and what norms logic imposes on either, for any particular construal of reasoning. I will return to the questions of what reasoning is and what norms logic imposes on reasoning in Part 4 of this thesis.

In this part of the thesis (Part 3), I analyze in detail Steinberger's recent work on the normativity of logic for *belief*, as well as the relevant work on which it is based. This yields a set of plausible norms that logic imposes on belief. I then discuss what these norms say regarding some particularly troublesome cases, and conclude that the norms handle the troublesome cases successfully. Next, I discuss a potential problem for Steinberger's preferred norm: it builds what some epistemologists would regard as prudential considerations into a putatively epistemic norm. I attempt to provide a solution to this problem. I conclude that logic is normative for belief, in

³³ The classic papers in this area are Harman 1986, Sainsbury 2002, MacFarlane 2004, Streumer 2007, Field 2009, and Milne 2009.

³⁴ Steinberger 2019a and Steinberger 2019b.

that there is a set of plausible bridge principles linking facts about logic to normative facts about belief.

Three Kinds of Norms

As I just stated, the rest of part 3 will investigate the norms of logic for *belief*. These will presumably take the form of principles that link facts about logic to normative facts about belief. What kinds of principles are these? Florian Steinberger has argued that the debate around the normativity of logic has been muddled because the proponents and opponents of logic's normativity have been discussing different kinds of normative principles. Steinberger distinguishes between three kinds of normative principles that could hold between logic and belief:³⁵

Directives: First-personal principles which guide a subject in what they ought to do, choose, or believe.

Evaluations: third-personal principles which set standards or ideals against which to assess acts or states as good or bad, correct or incorrect, and so on.

Appraisals: third-personal principles which lay the basis of our attributions of praise or blame to agents.

Directives are *first-personal* insofar as they are the kinds of norms that are supposed to offer guidance to an agent; they answer the question, “what should *I* do?” for an agent in a given situation. Evaluations and appraisals are *third-personal* insofar as they are used to

³⁵ Steinberger 2019a, p. 16.

evaluate/appraise other agents. According to Steinberger, while what Harman sought was a directive, those who have attempted to defend the normativity of logic have typically offered evaluations or appraisals.³⁶

Here's a mundane example that may help make sense of the different kinds of norms. Suppose you are driving to work, and your favorite route is blocked-off by construction. You employ directives in order to assess which alternate route to take, say, route A instead of route B. Perhaps you think, "I should take the more direct route, as it is likely to be faster. Route A is more direct, so I'll choose that." The 'should' in the first sentence has the force of a directive. By contrast, other people (perhaps your boss!) employ evaluations to assess whether you ultimately made the right decision. When you finally arrive late for work and explain what happened, your boss might say, "you should have taken route B, as you would have gotten here 10 minutes earlier!". The 'should' in this sentence is an evaluative-should. Your boss may then employ appraisals to further ascribe blame or praise onto you, in light of what he thinks you are blameworthy for. Your boss may withhold blame if he thinks you could not have known that route B would be faster, or that there would be construction on your usual route. But if your boss instead thinks that you should have known better, he may blame (and perhaps reprimand) you for your decision.

This example highlights a few key points. First, a good directive must be guidance giving, and will be highly sensitive to an agent's limitations, while an evaluation sets an ultimate standard of correctness, and need not be sensitive to an agent's limitations. "Take the faster route!" is a fine evaluation for the situation in the last paragraph, but is a useless directive unless

³⁶ Steinberger 2019b, p. 318.

you already know which route will be faster. “Take the more direct route!”, by contrast, might be a good directive, provided that you are in a position to know which route is more direct, and that the more direct route is faster *ceteris paribus*. But it’s a bad evaluation, if your aim is to arrive at your destination quickly: if taking the more direct route causes you a significant delay, there is an important sense in which you made the wrong decision, despite you perhaps not being in a position to know that ahead of time. So evaluations may be highly idealized, while directives will generally not be.

Second, appraisals can vary in their degree of idealization between the degrees of evaluations and of directives. An appraisal which governs how blameworthy you are for an action will likely be quite sensitive to what you know or should know. But (somewhat controversially) there also may be appraisals which are less idealized. For example, a legal appraisal may judge you as legally blameworthy and hence deserving of reprimand whenever you fail to live up to the evaluative standard of the law – i.e., whenever you break the law – regardless of whether you knew you were breaking the law or were in a position to avoid doing so. Of course, such an appraisal may be unfair, but it serves as a good example of a legal appraisal which is insensitive to the limits of the lawbreaker nonetheless.

While my earlier examples were of mundane prudential norms and legal norms, the three different kinds of normative principle will exist in any domain of normativity, including the epistemic. These principles must be kept apart, though there may be close relations between norms of certain kinds. For example, perhaps there is a close link between directives and those appraisals which assess blameworthiness. But I won’t assume any such links here.

As a final note for this section, I’ll point out that the distinction between these three kinds of norms raises questions about how they should fit in to the picture of epistemic normativity I

gave in Part 2. Indeed, if we grant that there are distinct directive and evaluative senses of *ought*, *is permitted to*, and *has reason to*, then we will need to refine the picture I gave earlier, and assess which sort of permission/obligation corresponds to epistemic justification. I'll return to the question of how to fit this distinction into my earlier picture in Part 5, when I discuss how the norms of logic relate to other epistemic norms.

Bridge Principles

Returning to the normativity of logic, Steinberger's distinction lets us formulate three separate questions about the normativity of logic for belief.³⁷

Q1. Is logic normative in the sense of providing directives?

Q2. Is logic normative in the sense of providing evaluations?

Q3. Is logic normative in the sense of providing appraisals?

What kind of answers can we give to these questions? Following John MacFarlane, we might proceed by providing *bridge principles* between facts about logic and normative facts about belief.³⁸ Steinberger generalizes the bridge-principle schema as the following.³⁹

(BP) If $A_1, \dots, A_n \models C$, then $N(\text{if } a(A_1), \dots, a(A_n), \text{ then } b(C))$.

Where:

- A_1, \dots, A_n are premises,

³⁷ Steinberger 2019a, p. 7.

³⁸ MacFarlane 2004, p. 6.

³⁹ Steinberger 2019b, p. 312. I have slightly modified Steinberger's notation.

- C is a conclusion,
- \vDash is the entailment relation of the logic in question,⁴⁰
- N is a normative claim, consisting of a deontic operator (*ought, is permitted to, or has (defeasible) reason to*)
- a is an attitude (belief, disbelief, suspension) toward the premises, and
- b is an attitude toward the conclusion.
- For simplicity, both instances of ‘if... then...’ are to be construed as expressing the *material conditional* (though we might also ask about whether there are plausible principles which make use of other conditionals).⁴¹

We might also relativize the antecedent of the conditional to a doxastic attitude, or qualified doxastic attitude, such as belief, rational belief, knowledge, etc. The resulting restricted bridge principle schema would then have the form:

(RBP) If $c(A_1, \dots, A_n \vDash C)$, then $N(\text{if } a(A_1), \dots, a(A_n), \text{ then } b(C))$

Where c is a doxastic attitude toward the claim of logical implication between the premises and the conclusion, and the other symbols are as before.

We can generate many different putative bridge principles by modifying any of the parameters just mentioned. Importantly, we may also vary the scope of the deontic operator. The deontic operator may plausibly take three possible scopes:

⁴⁰ Presumably, genuine bridge principles would only hold between the norms of reasoning and whatever logic(s) correctly characterizes the validity of arguments. Anyone who holds a strong anti-realist stance toward whether any logic(s) characterize an underlying notion of validity is unlikely to find the search for a bridge principle between logic and reasoning very important.

⁴¹ Steinberger 2019b, p. 312.

(C) $P \rightarrow N(Q)$ (the operator only scopes over the conclusion)

(B) $N(P) \rightarrow N(Q)$ (the operator scopes over both the premises and conclusion, individually)

(W) $N(P \rightarrow Q)$ (the operator scopes over the entire conditional.)

All of the preceding discussion has been kept at maximal generality, but it will help to see some examples. Here are a few example bridge principles which will be mentioned later:

(CO+) If $A_1, \dots, A_n \models C$, then if S believes each of A_1, \dots, A_n , S ought to believe C

(WO+) If $A_1, \dots, A_n \models C$, then S ought to (believe C, if S believes A_1, \dots, A_n)

(WO-) If $A_1, \dots, A_n \models C$, then S ought (not disbelieve C, if S believes A_1, \dots, A_n)

(WR+) If $A_1, \dots, A_n \models C$, then S has reason to (believe C, if S believes A_1, \dots, A_n)

(WR+B) If S believes ($A_1, \dots, A_n \models C$), then S has reason to (believe C, if S believes A_1, \dots, A_n)

Before moving on, it will be worthwhile to give a brief discussion about the natural language readings of the various principles. In presenting the wide-scope principles (the W's), Steinberger uses parentheses to delineate the scope of the normative operator in question. Giving adequate natural-language glosses of these principles is tricky, but the following locutions are plausible readings:⁴²

(WO+) If $A_1, \dots, A_n \models C$, then S ought to see to it that, if they believe each of A_1, \dots, A_n , they believe C.

⁴² This is the gloss used by MacFarlane. See MacFarlane 2004, p. 7.

(WO-) If $A_1, \dots, A_n \models C$, then S ought see to it that, if they believe each of A_1, \dots, A_n , they do not disbelieve C .

(WR+) If $A_1, \dots, A_n \models C$, then S has reason to see to it that, if they believe each of A_1, \dots, A_n , they believe C .

(WR+B) If S believes $(A_1, \dots, A_n \models C)$, then S has reason to see to it that, if they believe each of A_1, \dots, A_n , they believe C .

Further, since the inner conditional is being taken as a material conditional, we can also give the following equivalent and unambiguous glosses of each principle using well known abbreviations of the material conditional in terms of ‘not’, ‘or’, and ‘and’. Here are some translated principles equivalent to WO+ (similar glosses can be given for the other principles):

- a. (WO+) If $A_1, \dots, A_n \models C$, then S ought to either not believe all of A_1, \dots, A_n , or believe C .
- b. (WO+) If $A_1, \dots, A_n \models C$, then S ought to not both believe each of A_1, \dots, A_n , and fail to believe C .

So there are adequate natural-language translations for the principles Steinberger discusses. However, in what follows, I will follow Steinberger in explicitly demarcating scope using parentheses.

It’s easy to see that varying the parameters will result in a plethora of possible bridge principles.⁴³ How are we to decide between them? Steinberger, MacFarlane, Harman, and others have each made numerous critiques against various bridge principles. From these critiques,

⁴³ The exact number of available principles depends on how many options we consider for each parameter.

general desiderata that any plausible bridge principle should satisfy can be established. In the next section, I will discuss each of these in depth.

Desiderata for Bridge Principles

Steinberger, drawing on MacFarlane and others, provides a total of eight different desiderata for bridge principles to satisfy. These are the following (descriptions of each consideration are my own, though I draw heavily from Steinberger):

1. **Belief Revision:** Whatever the link between facts about entailment/deduction and facts about reasoning, facts about entailment do not straightforwardly normatively compel one toward any particular attitude toward a conclusion. As we saw before, believing P and $P \rightarrow Q$ does not compel me to believe Q , as it may be better for me to instead reject P or $P \rightarrow Q$. Whatever normative principles link logic and belief, they must be flexible enough to accommodate the different acceptable responses to facts about entailment.⁴⁴

Steinberger attributes to John Broome a special case of Belief Revision, which deserves its own mention.⁴⁵

2. **Bootstrapping:** Since every proposition trivially entails itself, a poorly chosen bridge principle risks a form of epistemic bootstrapping. Suppose our principle were $CO+$ above. Then, since $P \models P$, our principle says that, if S believes P , then S ought to believe P . This is clearly unacceptable.⁴⁶

As both Steinberger and MacFarlane note, the first two considerations, Belief Revision and Bootstrapping, tell us that acceptable bridge principles must have their normative operator range

⁴⁴ This is discussed in Steinberger 2019a, p. 11.

⁴⁵ See Steinberger 2019b, p. 310, and for the original source, Broome 2000, p. 85.

⁴⁶ This is discussed in Steinberger 2019b, p. 310.

over both the antecedent and the conclusion of the inner conditional. This tells against narrow scoped principles like CO+.⁴⁷

3. **Excessive Demands:** Humans are cognitively finite beings. It seems unreasonable for the norms of logic to demand more than any human could possibly achieve. If the correct normative principles demanded, for example, that our beliefs be closed under logical implication, this would violate an “ought implies can” principle. To avoid this, any acceptable bridge principle must thus be sensitive to the agent’s cognitive limitations, and not impose excessive demands.⁴⁸

4. **The Priority Question:** As MacFarlane says, “we seek logical knowledge so that we will know how we ought to revise our beliefs; not just how we will be obligated to revise them once we acquire this logical knowledge, but how we are obligated to revise them even now, in our current state of ignorance.”⁴⁹ That is, it seems that the norms of logic should not only apply to us when we recognize them, but rather, whether we recognize them or not. This seems to indicate that any principle which imposes weaker constraints on an agent’s beliefs the less logical knowledge the agent possesses has got things backward.

Excessive Demands and the Priority Question pull in opposite directions. Excessive demands points us toward attitudinally restricted principles, or at least toward principles which are weaker than WO+. The Priority Question, on the other hand, points away from attitudinally restricted principles, as it imposes demands on agents who don’t have knowledge or beliefs about the logical entailments in question.

5. **Clutter Avoidance:** As Harman noted, any proposition you believe will trivially entail an infinite number of propositions, many of which are entirely unimportant to you. For example, P entails $P \vee Q$, for any Q

⁴⁷ For narrow scoped principles that are less strong than positive claims about obligation, Belief Revision and Bootstrapping may be less decisive. However, as we will see, these approaches are independently less well-motivated than the wide-scoped principles.

⁴⁸ Steinberger 2019b, p. 311.

⁴⁹ Steinberger 2019a, p. 14.

whatsoever. As Steinberger puts it, “it would be positively irrational for [you] to squander [your] meager cognitive resources on inferring trivial implications of [your] beliefs that are of no relevance to [your] goals.”⁵⁰

Clutter Avoidance immediately engenders a worry: I said earlier that I will focus on the normativity of logic for *belief* throughout part 3, but Clutter Avoidance seems to be a norm on what you may *infer*. But note that, even if Clutter Avoidance as stated is a constraint on what you may infer, certain norms on belief may nevertheless violate it. Any principle which would provide unqualified obligations to believe the logical consequences of your beliefs, like WO+, will plausibly require you to *come to believe* these consequences, and this may require *inferring* those consequences.⁵¹ No correct principle of the normativity of logic could plausibly demand you spend cognitive resources deriving propositions that have no value to you.

6. The Preface Paradox: There are situations where it seems like one epistemically ought not have logically closed beliefs. The Preface Paradox is a particularly potent example. Presumably, you have good reasons to believe thousands of propositions. But you also know that you are fallible: as you are unlikely to be right about everything you believe, it seems you would be justified in believing that at least one of your beliefs is false. But if so, then while you would assent to everything you believe, you are justified in not assenting to the conjunction of everything you believe. This is the Preface Paradox, and seems to indicate that, at least in some cases, you should not believe what is entailed by your beliefs.⁵²

7. The Strictness Test: As Steinberger puts it, “At least when it comes to ordinary, readily recognizable logical implications leading to conclusions that the agent has reason to consider, the

⁵⁰ Steinberger 2019a, p. 11.

⁵¹ There are complex issues regarding the relationship between belief and reasoning that must be dealt with to fully address the worry, but I will discuss these in detail later on.

⁵² Steinberger 2019a, p. 12.

logical obligation should be strict: there is something amiss about an agent who endorses the premises and yet disbelieves the conclusion on account of stronger countervailing reasons.”⁵³

8. **Logical Obtuseness:** At least in ordinary cases, an agent who believes some premises but refuses to take a stand toward their obvious consequences seems to exhibit an epistemic failing. This is similar to the strictness test, though applies to conclusions we suspend judgement on, rather than those that we disbelieve.⁵⁴

The Strictness Test tells us that, at least in ordinary cases, our bridge principle must be quite strong, and tells against principles which use the reasons operator. Similarly, Logical Obtuseness tells us that we must accept or reject conclusions which we see follow from our beliefs: merely not disbelieving (suspending on) an obvious conclusion of premises we believe is not enough. These desiderata also seem to contradict the core lesson of the Preface Paradox, which tells us that, at least in some cases, we can be justified in not believing the consequences of our beliefs. There is much more to be said about the Preface Paradox, and I will return to it later.

The following table (Table 1) may make it easier to keep in mind exactly what each desideratum tells us:

⁵³ Steinberger 2019b, p. 315.

⁵⁴ Steinberger 2019b, p. 316 (footnote 16). See also MacFarlane 2004, p. 12.

Table 1: Summaries of Bridge-Principle Desiderata

Given the following desideratum:	a good bridge-principle should...
Belief Revision	be flexible enough to accommodate cases where one should reject the premises, and cases where one should accept the conclusion, in light of an entailment fact.
Bootstrapping	not permit epistemic bootstrapping: that you believe P should not entail you ought believe P.
Excessive Demands	not impose excessive demands on agents.
The Priority Question	impose requirements that apply to agents whether or not they are aware of the entailment facts.
Clutter Avoidance	not require an agent to infer useless propositions.
The Preface Paradox	not require logically closed beliefs in preface-like cases.
The Strictness Test	impose strict (not easily overruled) requirements.
Logical Obtuseness	prohibit agents from simply taking no stand on the logical entailments of their beliefs.

As we've noted, Excessive Demands is in tension with the Priority Question, and Strictness and Logical Obtuseness seem to conflict with the Preface Paradox. So it would appear that all of our bridge principles conflict with at least one of our desiderata: those that make it past the Belief Revision, Bootstrapping, and Clutter Avoidance criteria either run afoul of the Preface Paradox or the Strictness Test. Further, if they are attitudinally restricted, the principles will violate the Priority Question, and if they are not, they *might* violate Excessive Demands. One might worry at this point that we must either reject one or more of our desiderata, or else give up on finding a suitable bridge principle.

But not so fast. It is now that Steinberger's innovation comes into play. For as Steinberger noted, there are three kinds of normativity for which we may seek bridge principles: evaluations, directives, and appraisals. And crucially, because these kinds of normativity serve

different purposes, *they are subject to different desiderata*. Table 2 below indicates which desiderata apply to which norms.⁵⁵

Table 2: Bridge Principle Desiderata and the Norms to which they Apply

	Evaluations	Directives	Appraisals
Belief Revision	X	X	X
Bootstrapping	X	X	X
Excessive Demands		X	X*
Priority	X		X*
Clutter Avoidance		X	X*
The Preface Paradox		X!	X!*
Strictness	X	X!	X!*
Obtuseness	X	X!	X!*

Some explanation is in order. First, the occurrence of an ‘X’ indicates that the corresponding consideration applies to the kind of norm in question. The occurrence of ‘*’ indicates that the consideration applies relative to the degree of idealization at play. As was noted before, appraisals can vary in their degree of idealization. More idealized appraisals will be subject to the same desiderata as evaluations, while less idealized appraisals are like directives. The occurrence of ‘!’ indicates that the norm is context sensitive, in a way that will be discussed below.

⁵⁵ This table is my extension of Steinberger’s table, which did not explicitly discuss Bootstrapping or Logical Obtuseness. See Steinberger 2019a, p. 24.

It is worth briefly explaining the rationale Steinberger gives for why the desiderata apply to the norms they do. First, Belief Revision and Bootstrapping apply across the board: any norm, no matter how idealized, should be flexible enough to accommodate the varying appropriate responses to facts about entailment, as we saw before. The best response to these objections is to opt for wide-scoped principles.⁵⁶ By contrast, Clutter Avoidance is important precisely because of the cognitive limitations of the agent in question, and so does not apply to evaluations and highly idealized appraisals.⁵⁷

I noted above that Excessive Demands and the Priority Question were in tension, as they seemed to have opposite recommendations regarding attitudinally restricted bridge principles. But as can be seen, this tension is avoided once it is recognized that the two desiderata apply to different kinds of norms. Since Excessive Demands is relative to the capabilities of an agent, it naturally applies less the more idealized the agent's cognitive capabilities. It thus serves as an appropriate desideratum for directives, but not for evaluations. By contrast, Steinberger argues that the Priority Question does not apply to directives. MacFarlane rules out the attitudinally restricted bridge principles (principles that are relativized to when an agent has a particular attitude toward the entailment fact) on the ground that, according to those principles, "The more ignorant we are of what follows logically from what, the freer we are to believe whatever we please—however logically incoherent it is."⁵⁸ But while, from an evaluative standpoint, a more logically ignorant person is worse off, it would be unreasonable for a directive to demand that a person respect logical entailments of which they are unaware. As Steinberger puts it, "Directives... can only go on the logical entailments that we, in our best estimation, take to

⁵⁶ See Steinberger 2019a, p. 24.

⁵⁷ See Steinberger 2019a, p. 24.

⁵⁸ MacFarlane 2004, p. 12.

obtain.”⁵⁹ Of course, this is not to say that a directive should permit one to forget logical knowledge in order to be subject to less stringent norms, but attitudinally restricted bridge principles say nothing about that one way or another. And plausibly, there are independent epistemic and perhaps prudential norms which forbid one from forgetting logical knowledge. The upshot is that attitudinally restricted bridge principles make for bad evaluations, but are perfectly acceptable as directives.

We are left with the Preface Paradox, the Strictness Test, and Logical Obtuseness. As Steinberger puts it: “From our strict evaluative point of view an agent who professes to believe a number of premises while at the same time disbelieving a logical consequence thereof is plainly logically defective.”⁶⁰ The same would go for refusing to take a stand on a logical consequence of the premises. Steinberger argues that the Preface Paradox, unlike Strictness and Logical Obtuseness, does not apply to evaluations.⁶¹ Here’s a version of his argument: Let B be the set of your beliefs, and let E be the set of propositions which you, from an evaluative standpoint, ought to believe. Steinberger thinks that, plausibly, the beliefs in E will be either true or known propositions: that is, they will be factive. The Preface Paradox arises when, for some (large) subset b of B, you ought believe everything in b, but not believe the conjunction of the propositions in b. But from an evaluative standpoint, this does not happen: if everything in b is in E, you should believe everything in b, as well as the conjunction of everything in b (and the other beliefs in E). On the other hand, if something in b is not in E, you should not believe it, or any conjunction containing it. So, from an evaluative standpoint, the Preface Paradox does not

⁵⁹ Steinberger 2019a, p. 23.

⁶⁰ Steinberger 2019a, p. 26.

⁶¹ Steinberger’s actual argument is both very quick and quite complex. This is my best reconstruction. See Steinberger 2019a, p. 26-27.

obtain.⁶²

However, both the Preface Paradox and the Strictness Test (and Logical Obtuseness) do apply to directives. This seems to be a problem, as the Preface Paradox points in the direction of principles which use the *reasons* operator, while the Strictness Test points away from said principles. There are at least three solutions in the literature, so long as one wishes to focus on the traditional doxastic attitudes, rather than credences.⁶³ The first, following Peter Milne, is to simply deny the Preface Paradox, and say that you should, in fact, believe the conjunction of everything you rationally believe, despite knowledge of your own fallibility.⁶⁴ Personally, I find this extremely attractive. But denying the Preface Paradox is highly controversial, and I don't believe that the normativity of logic should hinge on it, so I will not argue for it here.⁶⁵

The second option is to follow MacFarlane, and to permit epistemic dilemmas.⁶⁶ MacFarlane thinks that, in Preface scenarios, one has two conflicting epistemic obligations: an obligation to believe consistently, and an obligation to recognize your own fallibility. All things considered, in Preface Scenarios, one should disbelieve the conjunction of the belief set in question, because of the overriding force of one's self-knowledge of fallibility. But, unless one knows which beliefs in the set are false, one is forbidden from revising any one in particular. So the agent is in an epistemic dilemma: no course of action is acceptable. MacFarlane uses this to

⁶² Note that this does not prescribe that an agent should (in the evaluative sense) have contradictory beliefs. The reason is that, from an evaluative standpoint, you should believe whatever subset of your preface beliefs A_1, \dots, A_n are true, and the conjunction thereof, $B_1 \& \dots \& B_n$. But despite your knowledge of your own fallibility, you (evaluatively) should not believe $\sim(B_1 \& \dots \& B_n)$, for the simple reason that every conjunct in that conjunction is true.

⁶³ The Preface Paradox can be easily dealt with if one focuses on credences. This is the approach taken in Field 2009.

⁶⁴ See Milne 2009, p. 285. In brief, Milne thinks that, as a norm of assertion, the size of the set of propositions in question never licenses you to assert each proposition in the set, but deny the conjunction thereof. Milne thinks this also extends to belief. Hence, no matter how large the set, if you are justified in believing each proposition individually, you are justified in believing the conjunction thereof.

⁶⁵ The main arguments against denying the Preface Paradox are to be found in Christensen 2004.

⁶⁶ See MacFarlane 2004, p. 13-14.

motivate adopting what I called WO- above, in conjunction with WR+, as the accepted bridge principles.

The third option, and the one I will endorse going forward, comes from Steinberger. Steinberger observes that the Strictness Test and the Preface Paradox seem at their most forceful in very different scenarios. If you should accept an atomic proposition A and another atomic proposition B, then it seems implausible that it could be reasonable for you to deny $A \wedge B$, even if you had countervailing epistemic reasons against believing $A \wedge B$. On the other hand, when the conjunction in question has, say, a million premises, that's when Preface Intuitions seem the most forceful. But these are very different kinds of cases, and there is a spectrum of more-to-less preface-like cases in between these two extremes. Steinberger's innovation is to note that the strength of a norm can vary with context: plausibly, in the case of a conjunction of atomic sentences the norm is decisively strong. But in the case of a million premise conjunction, it might be rather weak, and hence can be overridden by considerations like the acknowledgement of one's own fallibility. To account for this, Steinberger advocates for principles which use a *reasons* operator, but whose operator varies in strength depending on the context: in an ordinary context, the reasons will most likely be decisive, but as the context becomes more preface-like, the reasons get weaker and eventually may be overridden.

One might worry that this account does not make clear exactly when Strictness will give way to Preface concerns. Steinberger argues that this is unproblematic: while the account is vague, the vagueness in the account is due simply to the vagueness in the subject matter.⁶⁷

⁶⁷ See Steinberger 2019b, p. 323-324.

Whatever account of vagueness one thinks is correct, it should be possible to simply import that account into the interpretation of the reasons operator.

The Correct Bridge Principles

We are now in a position to look at the bridge principles which satisfy our desiderata for each normative role. First, directives. Steinberger, after accounting for all of the appropriate desiderata, comes to the following principle:

(WR+B*) “If S believes that $A_1, \dots, A_n \models C$ and S considers C or has subjective reasons to consider C, then S has reasons to (believe C, if S believes all of the A_i).”⁶⁸

This principle is synchronic, though we could use it to generate a corresponding diachronic principle if desired.⁶⁹ It also satisfies all of Steinberger’s desiderata for directives, notably ignoring the Priority Question, and thus being attitudinally restricted. The qualification, “if... S considers C or has subjective reasons to consider C” handles Clutter Avoidance worries, though this verbiage is noteworthy, and I will return to it. The principle is also wide-scope, which handles Belief Revision and Bootstrapping worries, and uses Steinberger’s contextually strong *reasons* operator, which handles both the Preface Paradox and the Strictness test. It’s also

⁶⁸ Steinberger 2019b, p. 321.

⁶⁹ Steinberger thinks this principle generates the following diachronic principle: “(†) If, at time t, S believes that $A_1, \dots, A_n \models C$ and S considers C or has subjective reasons to consider C, then if S’s reasons for believing the A_i are not outweighed by sufficiently strong prior reasons for doubting C, then S has reasons to believe C at t_0 (where t_0 is preceded by t); if S has sufficiently strong prior reasons for doubting C, S has reasons to revise her belief in some of the A_i at t_0 in a way that is consonant with the evidence and such that the remaining set of believed propositions—a proper subset of $\{A_1, \dots, A_n\}$ —no longer entails C.” See Steinberger 2019b, p. 322.

positive, thus avoiding Logical Obtuseness concerns. For all that's been said, this seems to be an adequate directive for the normativity of logic.

But we can do better. First, Steinberger notes that the language of WR+B* is over-intellectualized: few people have explicit beliefs about logical entailment, yet our bridge principle seems like it ought to apply to non-logicians. To counter this, Steinberger suggests the following formulation, which I'll call Steinberger's Directive:

(SD) "If according to S's best estimation at the time, S takes it to be the case that $A_1, \dots, A_n \models C$ and S has reasons to consider or considers C, then S has reasons to (believe C, if S believes all of the A_i)"⁷⁰

This formulation handles the previous worry so long as the mentioned taking-condition is not understood in overly intellectualist terms: lay people can take it to be the case that $A_1, \dots, A_n \models C$ without having any explicit beliefs about the subject, and without any training in formal logic.⁷¹

It also handles another worry. WR+B imposed a directive on agents who believed a certain entailment holds. But one might have worried that this wasn't strong enough: what if one has wildly unjustified logical beliefs? Steinberger's Directive handles this by restricting the principle to entailments that hold in an agent's *best estimation*.⁷² For a proposition to hold in an agent's best estimation presumably just is for that proposition to be supported by the totality of that agent's evidence about that proposition, i.e. for the proposition to be justified. While it is

⁷⁰ Steinberger 2019b, p. 325.

⁷¹ The taking-condition is the subject of extensive literature, and is most notably defended in Boghossian 2014. I'll discuss the taking-condition as a condition on reasoning in part 4.

⁷² Another option would be to restrict the principle to entailments that the agent realizes (see Sainsbury 2002, p. 4) or that the agent knows (see MacFarlane 2004). These attitudes are factive, and hence, would restrict the entailments in question to entailments that actually hold in whatever logic(s) govern(s) the validity of natural language arguments. Arguably, factive attitudes are less suitable for *directives*, since presumably, one should (from one's own standpoint) believe consistently with the entailments that one justifiably believes obtain, even if these entailments in fact don't obtain.

certainly possible for a proposition that holds in the agent's best estimation to be *false*, it presumably will not be possible for the proposition to be unsupported. So this formulation of Steinberger's Directive prevents the directive from extending to unjustified beliefs about entailment, but still permits justified but false beliefs about entailment to have normative force on the agent's doxastic attitudes.

So Steinberger's Directive seems to provide a plausible directive for the normativity of logic for reasoning and belief. What about evaluations? As can be seen in Table 2, evaluations are not subject to Clutter Avoidance, The Preface Paradox, or Excessive Demands, and unlike directives, *are* subject to the Priority Question. This means that any appropriate evaluation will be a much stronger principle than in the case of directives. The following principle, which we saw before, is the strongest possible evaluation which satisfies all of these desiderata:

(WO+) If $A_1, \dots, A_n \models C$, then S ought to (believe C, if S believes all of the a_i)

WO+ is effectively a closure requirement on belief. It states that an agent's beliefs should contain every consequence of the beliefs in the set. Notably, it doesn't pay any heed to whether the agent is or could be aware of said consequences, nor to the cognitive requirements of forming these beliefs. Since it is wide-scope, it handles Belief-Revision and Bootstrapping, and thus still permits one to remove beliefs that entail a conclusion which one independently ought to reject. It handles the Priority Question since it is not attitudinally restricted, and, since the operator is an ought, and the principle is positive, it handles both Strictness and Logical Obtuseness. Because this principle is intended as an evaluation, that's exactly what is desired.

What about appraisals? As has been noted, appraisals vary in strength depending on the degree of idealization being applied. Minimally, it seems that an appraisal corresponding to

Steinberger's Directive would be appropriate for assessing blameworthiness: you are prima facie blameworthy if you fail to hold beliefs which are required by the logical entailments that hold in your best estimation, that you consider. Of course, since this is a prima facie norm, it is possible that you have overriding reasons not to hold such beliefs (perhaps stemming from a preface situation that you recognize), and in such a case, your blameworthiness would be exculpated. At the other end, WO+ offers an evaluation of optimal performance with respect to the normativity of logic, and might be suitable for stating a measure of ultimate praiseworthiness. And of course, it is possible for there to be appraisals which fall in between: perhaps, in certain contexts, a lower bar for praiseworthiness is warranted, or the standards for blameworthiness may be more strict.

Logical Truth, Logical Falsehood, and the Norms of Logic

It is worth briefly discussing how the norms of logic defended in the last section (WO+ and SD) handle certain potentially deviant cases. One might worry about the consequences of these principles regarding two particular kinds of cases: the cases of logical truths and logical falsehoods. Specifically, what do our principles prescribe when:

1. $\emptyset \models C$
2. For some inconsistent set of premises A_1, \dots, A_n and a clearly objectionable conclusion C ,
 $A_1, \dots, A_n \models C$ (in any logic that validates *Ex Falso Quodlibet*)

Plugging each example into our aforementioned norms, we get the following:

- 1a. if $\emptyset \models C$, then S ought to (believe C, if S believes all of the $A_i \in \emptyset$)

1b. “If according to S’s best estimation at the time, S takes it to be the case that $\emptyset \models C$ and S has reasons to consider or considers C, then S has reasons to (believe C, if S believes all of the $A_i \in \emptyset$)”

2a. (For some inconsistent set of premises A_1, \dots, A_n and clearly objectionable conclusion C) if $A_1, \dots, A_n \models C$, then S ought to (believe C, if S believes all of the a_i)

2b. (For some inconsistent set of premises A_1, \dots, A_n and clearly objectionable conclusion C) If according to S’s best estimation at the time, S takes it to be the case that $A_1, \dots, A_n \models C$ and S has reasons to consider or considers C, then S has reasons to (believe C, if S believes all of the A_i)

1a and 1b apply WO+ and SD to the case where C follows from no premises. This may occur whenever C is a logical truth. But in these cases, the prescriptions of our norms are unproblematic. As I noted before, WO+ is a closure requirement, which is appropriate for an evaluation. Since it is vacuously true that S believes all of the elements of \emptyset , 1a prescribes belief in every logical truth. While this obviously would make for an unacceptable directive, as an ultimate standard of correctness (an evaluation), the demand for logical omniscience is unproblematic. 1b, by contrast, says merely that agents have reason to believe the logical truths which hold in their best estimation. Since these cases are maximally not preface-like (they have no premises at all), these reasons will be decisive, and thus have the force of an ought. So 1b requires agents to believe the logical truths that hold in their best estimation, which is a very plausible directive.

2a and 2b arise for logics that contain the Principle of Explosion, such as classical logic. In these logics, any conclusion whatsoever follows from any contradictory set of premises. If one of these logics is taken to correctly describe logical validity, then there will be a plethora of instances where a patently false conclusion follows from an inconsistent set of premises. If 2a or 2b require agents of the appropriate sort (ideal agents in the case of 2a, agents who recognize the entailment in the case of 2b) to believe patently false conclusions, this would be problematic. Indeed, proponents of logics which deny EFQ have sometimes given arguments to this effect to support their preferred logical theory.⁷³

But it is important to keep in mind 1) that the *ought* operator takes wide-scope over the conditional, and 2) that both conditionals are *material conditionals*. So 2a is an *evaluation* which in effect states that, if a set of premises is inconsistent, you ought to either not believe each of the premises or believe the patently false conclusion. Similarly, 2b is a *directive* which states that, if in your best estimation an inconsistent set of premises entails a patently false conclusion C, you have reason to either not believe each of the premises or believe C. Both of these principles are compatible with there being cases where one ought to believe a contradiction, but they don't entail that there will be such cases. For 2a in particular, if, as Steinberger argues, the propositions that our complete set of epistemic evaluations would prescribe an agent to believe are all true, and if there are no true contradictions, then the true epistemic evaluations (including 2a) will never say it is correct for an agent to believe a contradiction. Hence, it will not be correct to believe a patently false conclusion derived from a contradiction either. So as an evaluation, 2a is unproblematic.⁷⁴

⁷³ See Steinberger 2016 for a discussion of the plausibility of these normative arguments for non-explosive logics.

⁷⁴ Of course, if – as some dialethesists would claim – there are true contradictions, then 2a may prescribe believing certain true contradictions. But in that case, 2a would be no more problematic than dialethism itself.

At first glance, it might seem that 2b can be handled in a similar way. It seems plausible that an agent could not have decisive epistemic reasons to believe a set of premises the agent had good reason to suspect was inconsistent, let alone a blatantly false conclusion which follows from those premises. If this were right, then 2b, in conjunction with whatever other epistemic directives there are, would simply prescribe the agent to reject the inconsistent set of premises, either by identifying the least plausible among those which generate the inconsistency and rejecting that, or else just suspending judgement on every member of the set.

But there is a problem with the suggestions of the previous paragraph, at least if we don't deny the preface paradox entirely. The preface paradox presents a family of cases where an agent's all-things-considered epistemic reasons favor holding on to an inconsistent set of premises, despite the agent recognizing the inconsistency. If an agent has good reason to believe a large set of premises (say, the contents of their own well-researched book), the preface paradox indicates it may be rational for that agent to believe each of the premises individually. But it may simultaneously be rational for the agent to believe that the conjunction thereof is false, on account of the recognition of their own fallibility. Since the set of premises and the negation of their conjunction are jointly inconsistent, any conclusion whatsoever will follow from them.

It is here that it becomes relevant that, in light of the preface paradox, we opted for the *reasons* operator when stating Steinberger's Directive. 2b states that an agent S has (contextually strong) reasons to make the conditional 'if S believes A_1, \dots, A_n , then S believes C' true when $A_1, \dots, A_n \models C$ in S's best estimation. But in the case of a rationally believed set of inconsistent premises and a patently false conclusion, there will be decisively strong reasons to reject C. These will plausibly overrule any reasons to believe C that arise from the fact that the agent is

rational in believing A_1, \dots, A_n .⁷⁵ In light of this, 2b need not commit any agent to being obligated to believe a patently false conclusion simply for recognizing that it follows from their (all-things-considered rational) beliefs.

Clutter Avoidance and Pure Epistemic Norms

We've now seen plausible bridge-principles for the normativity of logic which provide both an evaluative standard for beliefs to live up to, and a directive for agents to follow when forming their beliefs. We've also seen how these principles bound the strongest and weakest appropriate appraisals one might give when assessing the praiseworthiness or blameworthiness of an agent's beliefs. As a reminder, these were Steinberger's Directive (SD), and WO+:

(SD) "If according to S's best estimation at the time, S takes it to be the case that $A_1, \dots, A_n \models C$ and S has reasons to consider or considers C, then S has reasons to (believe C, if S believes all of the A_i)"

(WO+) If $A_1, \dots, A_n \models C$, then S ought to (believe C, if S believes all of the A_i)

⁷⁵ There's a case to be made that, even if we had opted for the stronger directive that arises from replacing the reasons operator in SD with the ought operator, there would still be good reason to think this would not require agents to believe patently false conclusions that follow from their rationally held inconsistent beliefs. As Steinberger has noted, it is a hallmark of wide-scope obligations that they are not detachable: one cannot infer from $O(A \rightarrow B)$ and $O(A)$ that $O(B)$ (Steinberger 2016, p. 409). This claim about the non-detachability of wide-scope oughts is equivalent to rejection of the axiom K in deontic logic: Steinberger is in effect asserting that $O(A \rightarrow B)$ does not entail $(OA \rightarrow OB)$. Steinberger claims that most deontic logicians in fact *do* reject K (Steinberger 2016, p. 410). In particular then, Steinberger argues that the fact that the agent ought to believe A_1, \dots, A_n and ought to believe that A_1, \dots, A_n entails C does not obligate the agent to believe C. If Steinberger is correct, then even this stronger principle would be unproblematic. But since we are making use of a version of SD with the *reasons* operator, we need not rely on Steinberger's argument here.

As we saw, these principles satisfy each of the desiderata that we have mentioned thus far. Thus, these principles seem like strong candidates for very plausible norms that logic imposes on beliefs.

But there is a remaining worry about Steinberger's Directive. The solution to Clutter Avoidance in Steinberger's Directive potentially is incompatible with certain conceptions of the epistemic. In order to account for Clutter Avoidance, Steinberger's Directive includes the proviso, "if... S has reasons to consider or considers C...". *Considering* a proposition is an action. But there are many epistemologists who explicitly reject the view that there can be epistemic reasons for actions of any kind, even mental actions like *considering a proposition*.⁷⁶ Such epistemologists hold the following view, which I'll call Only Doxastic Attitudes:

Only Doxastic Attitudes: The epistemic is a domain of normativity concerned only with the epistemic status of doxastic attitudes (i.e. truth-directed mental states) at a given moment, given the epistemic reasons for and against each attitude. While one may have *prudential* reasons to act in ways that are conducive to future positive epistemic states, there are no *epistemic* reasons for or against actions, including mental actions like *considering a proposition*.⁷⁷

If such epistemologists are correct, then Steinberger's Directive seems to build prudential reasons (the reasons one has to consider a proposition) into a putatively epistemic norm. But this is problematic, as it seems to violate the autonomy of the epistemic as a domain of normativity.⁷⁸

⁷⁶ One might worry that this view is trivially false, because *forming a belief* or *revising a belief* are mental actions that no theory of the epistemic could do without. Epistemologists the sort of view I have in mind may respond by insisting that belief formation is not voluntary, so the mental processes of forming or discarding beliefs is a mere activity, not an action. This response is predicated on a form of doxastic involuntarism. I take no stand on the issue of doxastic involuntarism here. Even if one does accept that *forming a belief* is an action which is epistemically assessable, there may be other non-ad hoc ways of demarcating these actions from non-epistemically assessable ones, like *eating a sandwich* or (perhaps) *considering a proposition*.

⁷⁷ For advocates of this view, see Cohen 2016, p. 850-853, Berker 2018, p. 456-461, Feldman 2000, p. 687-691, and Feldman 2003, p. 47-48.

⁷⁸ Of course, there are epistemologists who deny that the epistemic is autonomous: arguably, both pragmatists and pragmatic encroachers seek to give theories where the epistemic status of a belief is dependent in part on prudential

Suppose, for example, that evidence constitutes one's epistemic reasons for any given proposition. Whether a proposition is on balance supported by the evidence should not depend on the prudential reasons you have to consider that proposition: No matter how good or bad it would be for you to consider the proposition in question, that shouldn't change whether your evidence supports it. And it isn't clear why this would change on any other plausible conception of epistemic reasons. So if Only Doxastic Attitudes is correct, then Steinberger's Directive as it is currently stated implausibly builds prudential requirements into an otherwise epistemic norm. It would seem then that defenders of Steinberger's Directive must either reject the autonomy of the epistemic, or reject Only Doxastic Attitudes. But either option seems like an unnecessary burden for the proponents of the normativity of logic to bear.

To avoid this consequence, we could try to modify Steinberger's Directive to avoid building in prudential requirements. But whatever modification we adopt, our new directive should still not violate Clutter Avoidance. Even a defender of Only Doxastic Attitudes should still object to there being any requirement to perform useless inferences. From the perspective of Only Doxastic Attitudes, epistemic norms on belief should not impose requirements on what you should infer. And (prudential) norms on what inferences you should perform will plausibly forbid you from squandering your cognitive resources.

Perhaps Steinberger's Directive could be adequately revised by simply removing the language, "has reasons to consider" from the bridge-principle. Why does Steinberger include it?

Steinberger says the following (*italics mine*):

factors. But both of these views are highly controversial, and it would be very surprising if proponents of the normativity of logic were forced to deny the autonomy of the epistemic to maintain their view.

“Notice that it would not be enough merely to add the clause ‘and S considers C’ because there may be cases in which an agent fails to consider logical consequences she has good reasons to consider. *For example, I might, out of intellectual dishonesty, fail to take into account a damning consequence of my philosophical position of which I am otherwise aware. Surely, in such a case, I have good epistemic reasons to consider it.*”⁷⁹

Steinberger’s thought seems to be this: it would be an epistemic failing, in the process of deliberating about whether P, to not take into account a relevant consideration that one is aware of about whether P. This seems right: forming a belief in an intellectually dishonest way, such that a relevant consideration is ignored, certainly seems epistemically improper. Any appropriate epistemic directive should require an agent to factor in all relevant considerations of which the agent is aware.

But note how ‘consider’ is being used by Steinberger: the hypothetical agent is *aware* of the consideration in question, but refuses to consider it in the sense that they don’t include it in their deliberative process.⁸⁰ So perhaps we can accommodate Steinberger’s concern by relativizing Steinberger’s Directive to those propositions one is aware of, rather than to those propositions one considers. We might thus suggest the following:

⁷⁹ Steinberger 2019b, p. 321.

⁸⁰ In what does this “awareness” consist? I cannot hope to do justice to this question here. Clearly, when one deliberates about whether C, or asks about whether C, one becomes aware of the question whether C. But one is also presumably aware of said question if one asked about it at some point in memory, or perhaps if one forms a belief about whether C through some deliberative or non-deliberative process. Perhaps this awareness can even arrive implicitly, when one deliberates about an issue that is relevant to whether C. I intend my proviso to apply to whatever the correct account is of cognitive awareness of C, or the question, whether C.

(SD*) If according to S's best estimation at the time, S takes it to be the case that $A_1, \dots, A_n \models C$ and S is aware of the question of whether C, then S has reasons to (believe C, if S believes all of the A_i).⁸¹

Unlike Steinberger's Directive, SD* is an unobjectionably epistemic norm even if Only Doxastic Attitudes is true, as it does not mention reasons to consider or act in any other way. It also does not require S to form attitudes toward propositions that trivially follow from S's beliefs, unless S is already aware of these propositions. It therefore accommodates Clutter Avoidance. It also does not permit S to be intellectually dishonest, since S would need to be aware of the question whether C in order to fail to consider it in the way that Steinberger wishes to preclude.

Conclusion to Part 3

We've now seen, in the form of WO+, Steinberger's Directive, and any appraisals that fall between them, plausible normative principles that facts about logical consequence impose on belief. These principles satisfy all of the desiderata that MacFarlane and Steinberger considered, and more besides. We also saw, in SD*, a plausible directive that logic may impose which is compatible with views such as Only Doxastic Attitudes. As such, Harman's skeptical challenge seems to have been answered: there are plausible epistemic evaluations, directives, and appraisals on belief that are imposed by the logical entailment relations between propositions. This would seem to vindicate the common intuitions that logic constrains what beliefs one may hold.

⁸¹ It's worth stressing that, while considering whether C (in Steinberger's narrow sense) is *sufficient* for being aware of the question of whether C, one can be aware of the question of whether C without considering C in that narrow sense.

This is not to say that there can be no further reasons to doubt the normativity of logic. But at this stage, the burden of proof has shifted onto the opponents of the normativity of logic: we would need a new argument that logic cannot impose normative requirements on belief. All of the arguments that Harman gave against logic's normativity were factored into the desiderata that were considered earlier, and each has been answered.

But there are still unanswered questions about the normativity of logic. For one, while the aforementioned norms explain how logic is normative for *belief*, it still isn't fully clear what this means for logic's normativity for *reasoning*. After all, Harman's original goal wasn't to oppose the normativity of logic for belief, but rather to argue that logic had no special importance to reasoning, construed as "reasoned change in view".

If we take reasoning to be identical to the rules of belief revision, then the norms we have defended so far, in particular SD and SD*, do seem to answer Harman's challenge. If you have contextually strong epistemic reasons to hold beliefs which are in accordance with the facts about entailment which hold in your best estimation, then this will greatly restrict the belief revisions which are permissible. But Harman's view of reasoning as belief revision is only one possible view. Further, the norms that we have thus far accepted only link facts about *logical entailment* to what one should believe. But one might have thought that there should be a more robust link between the proof-theoretic rules of deduction and the rules of reasoning.

The time has come to investigate the nature of reasoning, and to see whether there are any norms which directly link facts about logic to facts about acceptable reasoning.

Part 4: The Normativity of Logic for Reasoning

In this section, I discuss the normativity of logic for reasoning. I start by discussing the nature of reasoning and the kind of account of reasoning that I will adopt, as well as what an account of reasoning should accomplish. I then discuss how the rules of reasoning might fit in to an overall epistemological theory, which was discussed briefly in Part 2. Next, I provide a sketch of a more comprehensive theory of reasoning, and discuss the advantages of my theory. One advantage is that it vindicates the normativity of logic for reasoning: on my view, the rules of natural deduction provide constitutive norms for the deductive reasoning. I show how this view handles all of the desiderata thus far considered. I then discuss how my account might fit into an overall epistemological theory. While, in the absence of a comprehensive theory of how reasons combine, what I say can be at best programmatic, I show that there is hope that my account will fit nicely into an overall theory, and that my account syncs up nicely with Steinberger's Directive.

What is Reasoning?

Reasoning is a ubiquitous part of our mental lives. Philosophers engage in reasoning when trying to form arguments for or against philosophical theses. Mathematicians engage in reasoning when trying to formulate solutions to mathematical problems, or when trying to prove mathematical theorems. Scientists engage in reasoning when trying to interpret the results of scientific experiments. But besides these academic applications, laypeople engage in reasoning in all sorts of situations. You engage in reasoning when you deliberate about which grocery store to go to,

or what to eat for lunch, or how much to tip the waiter. Any adequate account of what reasoning should account for all of these sorts of cases, and many more besides.

If logic is normative for anything at all, it seems like it ought to be normative for reasoning. As I stated in Part 3, there is a very strong intuition among laypeople and philosophers alike that there is a tight link between logic and reasoning. Laypeople often use phrases like, “that’s not logical” and “that’s not good reasoning” interchangeably. And the most stalwart opponents of the normativity of logic, such as Gilbert Harman, explicitly target the putative normative link between logic and *reasoning*. But what is reasoning anyway?

While there are numerous accounts of reasoning in the philosophical literature, there is no consensus on what reasoning is. Nevertheless, there are some commonalities among the theories. On anyone’s view, reasoning minimally consists in transitions from mental states with semantic content to other mental states with semantic content.⁸² This is a necessary condition for something to be reasoning, but it is far from sufficient. Further, on any account, reasoning occurs in steps: primitive steps of reasoning will consist of whatever minimal mental transitions comprise larger units of reasoning. To reason from P to Q will involve applying some sequence of primitive steps starting from P, and transitioning (via the conclusions of each primitive step, whose content is itself another proposition) to Q. Finally, on most views, reasoning (at least of the sort that is employed in the cases above) is a *personal*, rather than *subpersonal*, phenomenon: when agent’s reason in the examples above, they deliberate explicitly and consciously on whatever it is they reason about.⁸³ There are two upshots to this. First, the sort of deliberate

⁸² At least, paradigmatic cases of reasoning are like this, though some authors may seek conceptions of reasoning that do not restrict it to mental states with explicitly semantic content.

⁸³ This conception of reasoning as a personal level phenomenon is expressed in Boghossian 2014, p. 2. One author who notably disagrees with this conception (and who uses his disagreement to argue against Boghossian’s account of inference) is Richard 2019. For more on the Personal/Subpersonal distinction, see Drayson 2014.

reasoning we are interested in is a kind of action, not merely something that happens to you. Second, while most views of reasoning will take the relation between the mental attitude that is reasoned-from and the attitude that is reasoned-to to be a *causal relation*, the former attitude should not be a *mere cause* of the latter.⁸⁴

There's a commonly accepted distinction between two kinds of reasoning: *theoretical reasoning* and *practical reasoning*. We might say that theoretical reasoning bears on what one should believe, while practical reasoning bears on what one should do.⁸⁵ By extension, it's typically accepted that epistemic reasons bear on theoretical reasoning, while practical (prudential and moral) reasons bear on practical reasoning. Regardless, I'll only be concerned with *theoretical reasoning* here.

The definition of reasoning that I am going to adopt is the one defended in Boghossian 2014. Boghossian gives the following definition of *inference*, which I will take to be synonymous with reasoning:⁸⁶

(Inferring): S's inferring from p to q is for S to judge q *because S takes* the (presumed) truth of p to provide support for q.

This definition of what it is to infer makes reference to what Boghossian calls the *taking condition*, which he defines as follows:⁸⁷

⁸⁴ This similarly holds for accounts of basing: while most accounts of basing are causal, the reasons that people base their beliefs on are not mere causes for the belief, but rather reasons for which one believes.

⁸⁵ Some authors take the rules of theoretical reasoning to concern which *beliefs* you should form, while the rules of practical reasoning might concern which *intentions* you should form. I won't take a stand on whether this is correct here.

⁸⁶ Boghossian 2014, pp. 4.

⁸⁷ Boghossian 2014, p. 5.

(Taking Condition): Inferring necessarily involves the thinker *taking* his premises to support his conclusion and drawing his conclusion *because* of that fact.

So to reason from P to Q is to form the belief that Q because you take the presumed truth of P (which you believe) to support Q. This notion of *taking* in turn makes P a *reason* to believe Q. It is worth pointing out that we cannot require that the sense in which a reasoner takes their premises to support their conclusion be via explicit belief, as this overintellectualizes reasoning.⁸⁸ Plausibly, I can reason from P to Q without having an explicit belief that Q in some sense follows from P. To solve this, Boghossian explicates reasoning as fundamentally a *rule-following* endeavor: when I reason from P to Q, I do so by applying rules of inference which I (at least tacitly) accept. I therefore *take* P to justify Q precisely in that, since I'm willing to apply a rule of inference to infer Q from P, I treat P as a reason to believe Q.

Both the taking condition, and the rule-following conception of inference that Boghossian advocates for in order to satisfy it, are highly controversial, as are all competitor accounts. The taking condition is defended in Boghossian 2014, Broome 2013, and Valaris 2017. Opponents of the taking condition include Wedgwood 2006, Hlobil 2014, McHugh and Way 2016, and Richard 2019. Independently of whether an author accepts the taking condition, rule-following accounts of reasoning are also controversial. They are defended in Boghossian 2014, Broome 2013, and McHugh and Way 2017. Opponents of such accounts include Wedgwood 2006 and Valaris 2017.

⁸⁸ This is discussed in Boghossian 2014, pp. 6-9.

I'll be following Boghossian in taking reasoning to be an essentially rule-following activity which, by extension, satisfies the taking condition. Insofar as we are concerned with theoretical reasoning, reasoning from P to Q will involve a series of mental transitions from beliefs in premises (in this case, P) to beliefs in conclusions (ultimately, Q) via applications of rules which serve as primitive steps of reasoning. In applying these rules, an agent therefore takes P to be a reason to believe Q. In turn, this allows reasoning to be construed as a deliberate action that one partakes in, as well as a personal (rather than subpersonal) level activity.

Now, in order for the account of reasoning expressed in the last paragraph to be complete, we would need to specify what it is to follow a rule. This is a notoriously difficult endeavor, and I cannot do justice to it here. One thing that is important, however, is that the correct specification of following a rule must not be defined in terms of reasoning, on pain of circularity in conjunction with the above account of reasoning. This seems to require that rule-following be explicated in terms of special kinds of dispositions, or perhaps for an account of rule-following which treats rule-following to be an intentional but sub-personal activity.⁸⁹ In what follows, I will assume that some adequate account of rule-following can be offered.

Before we can move on, there are two more issues worth addressing. First, some instances of reasoning are good or correct, while others are not. If, in performing mental arithmetic, I calculate that $2 + 3 = 6$, I have made an error. Plausibly, this error stems from failing to apply an acceptable rule of inference: in trying to perform the + operation, I have erroneously performed some other operation instead. So one way to reason badly is to fail to apply a rule which I intended to apply. But I can also reason badly even if I apply the rule I

⁸⁹ See Broome 2013, ch. 13 and Broome 2014 for the former sort of account of rule-following, and Boghossian 2014, pp. 15-16 for a sketch of the latter.

intend: applying a rule of reasoning that corresponds to *affirming the consequent* plausibly will fail to count as good reasoning.⁹⁰ Even when the rule which is applied is truth-preserving (or otherwise provides the support of the conclusion that is intended), it can still count as bad reasoning. Boghossian gives the following example:

“Consider someone who claims to infer Fermat’s Last Theorem (FLT) directly from the Peano axioms, without the benefit of any intervening deductions, or knowledge of Andrew Wiles’s proof of that theorem. No doubt such a person would be unjustified in performing such an inference, if he could somehow get himself to perform it.”⁹¹

Such cases show that even truth-preserving inferences may not be acceptable, unless the rules of inference can be justified in some way or other. I’ll thus take good reasoning to be the kind of reasoning that employs only acceptable rules of inference.

What makes certain rules of reasoning acceptable? Some steps of reasoning seem to be primitive, while others can be derived from primitive rules, and thereby justified. But what justifies my use of a particular *primitive* rule of inference? This is a difficult question, and I won’t attempt to settle the issue here. But there are a number of possibilities. First, there is a long tradition of philosophers who take primitive rules of reasoning to be justified by a special faculty of *rational insight*, though opponents of this view tend to find it unsatisfying and unexplanatory.⁹² Second, Boghossian has argued that certain primitive steps of reasoning may be *blindly justified* in the sense that understanding the meaning of certain terms may require acceptance of certain rules of inference involving those terms. Such an approach relies on a form of inferentialism about the meanings of the terms in question.⁹³ There may be other approaches,

⁹⁰ I take this example from McHugh and Way 2017, p. 182.

⁹¹ Boghossian 2014, p. 6

⁹² See Boghossian 2003 for criticism of this approach.

⁹³ For defense of this idea, see Boghossian 2003.

but for my purposes here, I will take for granted that some rules of inference may be primitively justified by some means or other.

Of course, what I've said so far is merely the barest sketch of a general account of reasoning. A fully fleshed-out account would say precisely what the rules of (good) reasoning are, and exactly how each is justified. But what I've said thus far will serve as a framework for understanding what I take reasoning to be.

How Does Reasoning Fit into a Theory of the Epistemic?

I've now given a general sketch of what reasoning is: reasoning is an activity of performing mental transitions from attitudes with semantic content to others, in which the transitions are taken by the agent to be justifying of those later attitudes, in virtue of the fact that the transitions are performed by following rules that the agent accepts. While this sketch is controversial, it's no more controversial than any other competing account in the literature: there is no generally accepted account of reasoning. With my conception of reasoning in mind, we might now ask how reasoning is supposed to fit into our overall epistemological theory. The following is a thesis we might entertain:

Identity: The rules governing correct reasoning just are the rules governing belief revision.

The Identity thesis is quite popular. It is advocated by Gilbert Harman and Bart Streumer, and is also adopted by Boghossian and McHugh and Way, among others.⁹⁴ As Harman puts it, on this view, reasoning is “reasoned change in view”. Any norms on reasoning are therefore also rules of belief revision, and vice versa. As Harman has argued, given Identity, the rules of natural

⁹⁴ See Harman 1986, Streumer 2007, Boghossian 2014, and McHugh and Way 2017.

deduction (or any other logical rules of deduction) cannot be rules of reasoning. The reason for this is simple: we have, as a rule of logic, the rule of modus ponens: from P and $P \rightarrow Q$ we may infer Q . But this can't be a rule of belief revision, for reasons we have seen already: it is not the case that from my belief in P and my belief in $P \rightarrow Q$ I am straightforwardly licensed to form a belief in Q . In some instances, I should rather reject one of my premises.

While it was briefly discussed in Part 2, it will be worth revisiting how the rules of belief revision fit into an epistemological theory. Given the sort of account of the epistemic I gave in Part 2, there are synchronic norms on belief, such as Evidentialism about the epistemic ought, as well as $WO+$ and SD . These norms will ground corresponding norms of belief revision, which require or permit the agent to come to believe exactly the propositions which the synchronic norms demand. But there may also be further norms demanding that beliefs be formed or revised in certain ways. These will be more substantive norms on belief revision which do not correspond to the synchronic norms on *what to believe*. However, even these norms will have synchronic counterparts in the norms on the proper ways of *basing* beliefs. Taking an example from the possible norms on belief revision for *credences*, if there is a norm of credence revision which states that credences must be updated by conditionalization, then there will be a corresponding synchronic norm stating that for one to have *properly based* one's credences on one's evidence, one must have updated on that evidence by conditionalization. In other words, at the synchronic level, this sort of diachronic norm on belief revision will manifest as a restriction on what counts as *proper basing*.

If Harman and those who agree with him are correct, then the rules of reasoning just are the rules of belief revision, so what I just said regarding the rules of belief revision will ipso facto apply to the rules of reasoning. But plausibly, Identity as stated is false. The reason for this

is that reasoning does not seem required for all forms of belief revision. When I see my girlfriend's car pull into the driveway, I am plausibly required to form the belief that my girlfriend's car is in the driveway, and hence to add this belief to my set of beliefs (perhaps removing other beliefs in the process, most saliently, the belief that my girlfriend's car is not in the driveway, should I have had it). But no inference is required for this to occur: the requirement to perform a certain kind of belief revision in response to my perception does not force me to perform any reasoning whatsoever.⁹⁵ Of course, I may further *infer* from my belief that my girlfriend's car is in the driveway to the belief that my girlfriend is home from work. But though my perceptual belief about my girlfriend's car must satisfy the norms of belief revision, I may form it without engaging in any reasoning whatsoever.

So Identity, in its pure form, is false. But perhaps this is unfair to the defender of identity: what they meant was really that, when it comes to performing inferences, the rules which characterize acceptable inference are identical to the rules which characterize which instances of belief revision based on inference are acceptable. In other words, the rules of reasoning establish the acceptable ways of properly basing inferentially formed beliefs. So we might entertain the following view:

⁹⁵ There may be a caveat to this. It isn't clear to me whether *discarding* beliefs in light of contradictory information requires a step of inference highlighting the contradiction. On one view, when I see my girlfriend's car in the driveway, I form a belief that it is in the driveway. I'd then perform some reasoning to deduce a contradiction from my antecedent belief that the driveway was empty, and then discard the belief that the driveway is empty. On this view then, the discarding of my prior belief would require inference. On an alternate view, when I see my girlfriend's car in the driveway, I automatically discard my prior belief that the driveway is empty, without ever performing inference. It seems to me to be an empirical question which of these views is correct, and I take no stand on the answer here. My broader point is simply that the process of *adding* the belief "my girlfriend's car is in the driveway" does not require inference.

Identity*: The rules governing correct reasoning are a proper subset of the rules governing belief revision: specifically, they are the rules governing the revision of beliefs *through inference*.

But even Identity* faces apparent difficulties. The most serious problem is that the identification of the rules of (theoretical) reasoning with the rules of *belief* revision doesn't seem to explain the apparent constraints on *hypothetical reasoning*. Consider the following example from McHugh and Way:⁹⁶

“You might reason as follows:

Suppose Jane had a beer. Then there are none left. But then Jeremy will be upset. So, if Jane had a beer, then Jeremy will be upset.”

This example seems undeniably to be a case of reasoning, and further seems to be an example of perfectly good reasoning, provided certain details are filled in (like that you believe that there is only one beer left, and that Jeremy will be upset if he can't have another beer, etc.). Since each of these steps of reasoning is acceptable, they plausibly are licensed by whatever rules of reasoning there are.

But consider the second-to-last step of the example in question. What attitude do you take to the proposition “Jeremy will be upset”? You don't *believe it*, at least if you are taking Jane's drinking of a beer to be purely hypothetical. And while you plausibly *suppose* that Jane had a beer, it seems unnatural to say that you also suppose that Jeremy will be upset: you discover by reasoning that ‘Jeremy will be upset’ follows from your suppositions; you don't suppose it at the outset. Further, while you can permissibly suppose anything at all, you aren't licensed, at the step

⁹⁶ McHugh and Way 2017, p. 187.

of reasoning in question, to infer from your assumptions, say, that the moon is made of cheese. So the conclusion that you draw from your suppositions can't itself be a supposition. Hence, your attitude toward "Jeremy will be upset" can't be belief, and it can't be supposition. So what is it?

The upshot is that, in light of such examples, it appears that Identity* is false: the rules of theoretical reasoning cannot solely be concerned with the revision of *beliefs*. And even if we extend theoretical reasoning and Identity* to handle hypothetical reasoning and assumptions, it is still left unclear what attitude you take (or ought to take) toward the conclusions of steps of hypothetical reasoning.

There is, however, still a further problem with Identity*. The problem is that it seems plausible that one may engage in correct reasoning, and thus reason in accordance with the rules of reasoning, while nevertheless violating the rules of belief revision. For example, after finding that your roommate has not come home until morning multiple nights in the past few weeks (with no explanation), and knowing that he has been having a strained relationship with his significant other, you might infer through a perfectly acceptable chain of (defeasible) inference that your roommate is cheating on his significant other. But suppose that, despite your reasoning, and without having stronger countervailing reasons, you still refuse to believe that your roommate is cheating on his significant other (perhaps out of a desire not to believe your roommate would act unethically). In such a case, you'd surely be epistemically irrational: your beliefs would not be appropriately responsive to your epistemic reasons, and you'd be violating the rules of belief revision. But intuitively, you wouldn't thereby be violating any rules of reasoning. But if the rules of reasoning just were rules of (inferential) belief revision, it wouldn't be clear how it could even be possible to perform steps of reasoning without forming the

appropriate belief, let alone how you could do so without violating any rules of reasoning. So it would seem that both the activity of belief revision, and the norms that govern it, must be distinct from the activity and norms of reasoning.

So an adequate theory of reasoning should account for the potential for norms on hypothetical reasoning, and should specify what attitude I have, and ought to have, to each step in my reasoning from my premise beliefs (and suppositions) to my conclusion. It should also explain how it is possible for me to reason acceptably but still fail to revise my beliefs appropriately. Further, a theory of reasoning should recognize that reasoning comes in different kinds, and may offer different degrees of support to the conclusions that are drawn in relation to the premises. Some reasoning (namely, deductive reasoning) is meant to guarantee the truth of the conclusion given the truth of the premises. But not all reasoning need be like this: some reasoning may only be intended to give defeasible or probabilistic support to the conclusion.

With these desiderata in mind, I'll now sketch out a general theory of reasoning, which, if correct, will vindicate the intuitions that there is a tight link between the rules of reasoning and the rules of deductive inference. Despite the arguments I've just given, I'm going to argue for an account of reasoning which does vindicate a qualified version of Identity*. On my account, reasoning will in fact be a kind of belief revision, and the rules of reasoning will be a special subset of the rules of belief revision. Also, on my account (and contra Harman), the deduction rules of logic that we accept will provide normative constraints on the rules of reasoning.

A Sketch of a Theory of Reasoning

The theory of reasoning that I have in mind takes great inspiration from Pavel Tichý's views on the nature of *deductive argument*.⁹⁷ So it will be worthwhile to briefly discuss Tichý's views on the matter.⁹⁸

Tichý distinguishes between two different ways of construing the nature of deductive proof, which he calls the *one-dimensional view of inference* and the *two-dimensional view of inference* respectively. On the one-dimensional view, the rules of deduction operate on *propositions* (or the sentences which express them). They license a transition from some propositions to another proposition (sentence): An application of modus ponens consists in applying said rule to P and $P \rightarrow Q$ to yield Q . On the two-dimensional view, the rules of deduction operate on (single conclusion) *sequents*, defined as *antecedent/consequent* pairs. On this view, an application of modus ponens allows me to derive, from the sequents $\Gamma \vdash P$ and $\Sigma \vdash P \rightarrow Q$, the sequent $\Gamma \cup \Sigma \vdash Q$ (where Γ, Σ are sets of antecedents: premises and suppositions in the context of the inference).⁹⁹ The core difference, then, between the one-dimensional and two-dimensional views is that the two-dimensional view treats the premises and suppositions which are active in a context of inference to be ineliminable parts of the inferences that are made. Thus, while on the one-dimensional view, you derive a proposition from some premises, on the two-dimensional view, what you derive is rather an *entailment statement* (to use Tichý's term) linking the premises to the conclusion.¹⁰⁰

⁹⁷ Tichý views the rules of deductive argument as doubling as rules of reasoning. But his view is most directly applicable as a view on the nature of deductive proof, so I will present it as such.

⁹⁸ The views I discuss in the next paragraph are discussed in Tichý 1988, ch. 13. See also Pezlar 2013 for a succinct discussion of Tichý's views and core arguments on the matter.

⁹⁹ The specific formatting of the sequent calculus I offer here is given in Forbes 1994, pp. 131-133.

¹⁰⁰ Tichý 1988, p. 236.

Tichý gives two main arguments in support of the two-dimensional view.¹⁰¹ First, Tichý argues that the one-dimensional view misleadingly presents proofs as giving support for their conclusions. But this is clearly not the case: if I reason from hypothetical assumptions to a conclusion, this need not support the conclusion whatsoever. Second, Tichý notes that the purpose of giving proofs is ostensibly to *learn* something. But it is hard to make sense of how this could be true in the case of hypothetical reasoning on the one-dimensional view. To use Tichý's example, suppose, from the following purely hypothetical assumption:¹⁰²

Peter and Paul are spies

I infer the conclusion:

Peter is a spy

Ostensibly, I ought to be able to learn something from this inference: otherwise, what is the point of performing it at all? But I certainly don't learn the conclusion, because the conclusion only follows from my earlier hypothetical assumption. What I learn is rather that the conclusion *follows* from the premise: I learn the entailment statement that Peter and Paul are spies \vdash Peter is a spy.

In my view, Tichý's insight is sorely underappreciated.¹⁰³ But his insight doesn't just extend to how we ought to conceive of proof, but rather, to how we ought to conceive of *reasoning*. So here's the general account of reasoning that I have in mind. When we engage in

¹⁰¹ There is actually a third reason why Tichý espouses the two-dimensional view: Tichý seems to hold the Fregean view that valid inferences may only be derived from *truths* (see Tichý 1988, p. 236, and Pezlar 2013, p. 55). As I do not share this predilection, I will not discuss Tichý's arguments based on it here.

¹⁰² Tichý 1988, pp. 219-220.

¹⁰³ Tichý credits Frege (1906, p. 425 and 1979, p. 335) with the original insight about the two-dimensional nature of proof, and credits Gentzen (1934, 1936) with first developing a formalism which correctly represents it (the sequent calculus).

reasoning, we seek to establish support relations between the claims that we believe or are taking for granted in the context of our reasoning (either for a later reductio, or merely to derive hypothetical results) and the conclusions we derive. The nature of these support relations depends on the kind of reasoning we take ourselves to be performing: it might be deductive entailment, but it also might be a defeasible support relation. Minimally, we can take our arguments to be *deductive* or *inductive*, but I won't rule out the possibility that various sorts of "inductive" support relations might be of fundamentally different kinds (for example, inference to the best explanation may be a fundamentally different kind of reasoning than inductive generalization). While philosophers and mathematicians may have sophisticated theories of what differentiates deductive and inductive inference, even laypeople with no such training can recognize when their reasoning is supposed to *guarantee* the conclusion in light of the premises, or when it merely *makes the conclusion likely*.¹⁰⁴

It's important to note that, on the view I'm espousing, the content of an entailment statement (or more broadly, a *support statement*) is itself propositional. What is its content?

There are three possible contenders for the case where the reasoning is deductive:

- (1) $A_1, \dots, A_n \vdash C$ (C is derivable from A_1, \dots, A_n)
- (2) $A_1, \dots, A_n \models C$ (C is deductively entailed by A_1, \dots, A_n)
- (3) $A_1, \dots, A_n \rightarrow C$ (If A_1, \dots, A_n hold, then so does C)

While there are considerations that favor both (1) and (2), (3) can safely be discarded. There are two reasons for this. First, in engaging in hypothetical reasoning, if you infer Q from the assumed P, it is an additional step of reasoning to discharge the assumption P and infer $P \rightarrow Q$.

¹⁰⁴ Similar comments can be found in Boghossian 2014, p. 5.

While agents will sometimes elide this step, they often state it explicitly. But they don't take themselves to be merely reiterating that P entails Q: rather, they are performing a genuine step of inference with the conditional. Second, in forming the belief that P entails Q, one doesn't necessarily ever form the belief that $P \rightarrow Q$. You might not do so if your assumption was for another purpose, such as for a *reductio ad absurdum*. In fact, it's at least logically possible that an agent could come to understand the concept of entailment without coming to understand the concept of the conditional. Suppose there were such an agent, and the agent understood disjunction, and was reasoning following the rule of \vee Elimination. The agent might assume P, then reason to R, then assume Q, and reason to R, and then come to form the belief that $P \vee Q$ entails R, using \vee Elimination. Such an agent would at no point need to form a belief in the conditional $P \rightarrow R$, nor could they, so we've supposed. It thus makes more sense to construe the content of a support statement as that of (1) or (2) rather than (3).

The issue is trickier in adjudicating between (1) and (2). Technically speaking, deriving a conclusion C from some premises/assumptions A_1, \dots, A_n establishes the *derivability* of C from said premises, using a certain set of inference rules. It is a further step, which depends on the soundness of the deductive system, to establish that A_1, \dots, A_n entails C.

Nevertheless, it strikes me as more plausible that the content of a support statement is better represented as a fact about entailment, rather than derivability. First, while establishing derivability may not thereby yield entailment if the deductive system is not sound, it seems highly implausible that the deductive system of the logic that underpins natural language (or the metaphysical structure of the world) would be unsound. Second, laypeople don't understand the difference between derivability and entailment, yet are perfectly capable of holding the belief that a conclusion *follows* from a set of premises. And the way they will come to that conclusion

is by reasoning to it from those premises. So it strikes me as a better representation of the practice of reasoning to take reasoning to establish entailment facts (of the form of (2)) rather than derivability facts (of the form of (1)). Third, laypeople and logicians alike can seek to establish that a purported entailment fact *doesn't* hold (i.e., that $A_1, \dots, A_n \not\models C$) by giving a counterexample. But this only can show a flaw in an argument if it is granted that the argument was meant to show a fact about entailment, rather than mere derivability. So even though, on my account, reasoning will typically proceed by applications of deductive inference rules, I will take the content of (deductive) support statements to be that $A_1, \dots, A_n \models C$.

So on my view, when you engage in deductive reasoning, you form beliefs that $A_1, \dots, A_n \models C$. You can also believe (letting ' \Rightarrow ' stand for defeasible support/non-monotonic entailment) that $A_1, \dots, A_n \Rightarrow C$. That is to say, when you reason using defeasible methods (such as inference to the best explanation), plausibly the conclusion you draw is that a set of antecedents defeasibly supports a conclusion, though, as in the case of deductive reasoning, you don't necessarily take any particular attitude in the conclusion C thereby.¹⁰⁵ So on the view I'm espousing, the rules of reasoning are rules one may apply for forming beliefs whose content is of the form $A_1, \dots, A_n \models C$ or $A_1, \dots, A_n \Rightarrow C$, depending on the kind of reasoning at issue. Your reasoning is in turn good or correct insofar as it correctly applies rules which you are justified in using for that kind of reasoning.

Which rules are these? In large part, this will hinge on the sort of theory of logical entailment we accept (of the sort that I discussed in part 1). I'll discuss the available options shortly. But for now, we can make a more general remark: note that, in attempting to establish

¹⁰⁵ While I can't discuss the rules of defeasible reasoning here, I envision a system of defeasible inference rules of a similar kind to that of the deductive inference rules. Such a system could perhaps be offered by the derivation system of a non-monotonic logic, like that of Horty 2012.

that some premises/assumptions conclusively or defeasibly support a particular conclusion, I seem to be beholden to the standards of rigor required by the kind of reasoning I am employing. For example, if I take my reasoning to be deductive, then in engaging in said reasoning, I become beholden to whatever standards there are for deductive reasoning. In other words, by engaging in reasoning of a certain kind, I plausibly must satisfy the *constitutive norms* of the kind of reasoning I employ.

The notion of a constitutive norm traces back to Kant, but the notion I employ here is the one discussed by Steinberger, which he in turn draws from Searle.¹⁰⁶ Constitutive norms are norms which, “create or define new forms of behavior. The rules of football or chess, for example, do not merely regulate playing football or chess but as it were they create the very possibility of playing such games.”¹⁰⁷ Steinberger elaborates on the notion nicely:¹⁰⁸

“I cannot... opt out of conforming to the rules of chess *while continuing to count as playing chess*; in systematically violating the rules of chess and persisting in doing so even in the face of criticism, I forfeit my right to the title ‘chess player’. Unless one appropriately acknowledges that one’s moves are subject to the rules of chess, one’s activity does not qualify as playing chess.”

Now, it’s important to emphasize that it *is* possible to violate a constitutive norm. I can, for example, break a rule while playing hockey, without thereby ceasing to play hockey. But plausibly, whether I am playing hockey while breaking its rules or simply not playing depends on whether I take myself to be subject to the rules of the game. If, as a result of breaking a rule, I allow myself to be criticized (or penalized) according to the rules of hockey, then I’m still playing hockey. But if I ignore such rules, then I am simply no longer playing.

¹⁰⁶ See Steinberger 2017a, p. 149, and for the original source, Searle 1969, p. 33-34.

¹⁰⁷ Searle 1969, p. 33-34.

¹⁰⁸ Steinberger 2017a, p. 150.

So when I suggest that there are constitutive norms for certain kinds of reasoning, I mean that, in order for me to count as engaging in, say, deductive reasoning, I must take myself to be subject to the norms of deductive reasoning.¹⁰⁹ If I am reasoning deductively and apply a rule of reasoning that is not deductively valid, I reason badly. But I can only count as engaging in deductive reasoning if I take my arguments to be subject to a requirement to be deductively valid: if I purport to reason deductively, and you show me that my reasoning is invalid, and I treat this fact as insignificant, then I can't be said to be engaging in deductive reasoning.

It's a common objection to the claim that an activity has constitutive norms that said norms threaten to make it impossible to cheat at the activity or perform it badly. I've tried to answer this objection by emphasizing that you can, in fact, violate constitutive norms, but if you fail to take them to apply to you, you are engaging in a different activity. But one might object that, nevertheless, I could cheat at, say, soccer, and refuse to respond appropriately when caught if I thought it would somehow help me achieve my aim of winning the game. In such a case, I would not thereby cease to be playing soccer, despite my open refusal to abide by the rules.¹¹⁰

There are two things to say in response. First, the amount I can willingly deviate from the rules of a game while still counting as playing the game is a vague matter: uncontroversially, if I run around the soccer field tackling everyone, I'm not playing soccer, but other deliberate rule violations (including tackling the odd person) will not constitute ceasing to play soccer. Whether

¹⁰⁹ My view is quite similar to the view held by Frege of the rules of logic as establishing the "laws of thought" (see Steinberger 2017a for discussion). But unlike Frege, I'm restricting my claim to particular kinds of reasoning, rather than thinking in general. Further, I take the rules of (deductive) logic to only provide constitutive norms for the rules of deductive reasoning. Other kinds of reasoning will have their own sets of constitutive norms, which are perhaps embodied by whatever theory of non-monotonic logic that we ought to accept.

¹¹⁰ Thanks to Graeme Forbes for pressing this objection.

I'm abiding by the constitutive norms of an activity is a matter of degree, but that does not refute that there are such norms.

Second, even if the objection is fatal to the claim that some set of norms is constitutive of playing a certain game, this need not trouble my account of reasoning, for the simple reason that it isn't possible to cheat at reasoning. Deliberately violating the rules of reasoning would not get me closer to the goal of forming true beliefs, except by sheer accident. And while I certainly can reason badly, I still must be applying rules that I *take* to be justifying of the conclusions I draw in order to count as reasoning.

Let's take stock: on the view I'm suggesting, reasoning is a process of forming beliefs in *support statements* of various kinds through the application of rules which I take to secure the appropriate kind of support. The application of these rules is subject to constitutive norms of inference depending on the kind of inference that I take myself to be engaging in (i.e., the kind of support statements I take myself to be inferring). And in the case of *deductive reasoning*, I plausibly satisfy these constitutive norms by reasoning according to the rules of logic.

Which rules are these? In Part 1, I noted that there were two broad ways of characterizing logical consequence. One can take an *inferentialist* stance, according to which, to understand the meaning (or at least part of the meaning) of the logical constants just is to grasp their inference rules. Logical consequence, on this view, is defined in terms of derivability from the inference rules which characterize the logical constants. Alternatively, one can take a *denotationalist* stance, according to which, the meaning of the logical constants is set by the functions that they denote.¹¹¹ On this view, a statement entails another insofar as every interpretation of the non-

¹¹¹ I say "functions" rather than "truth-functions" to account for logical constants which are not straight-forwardly truth-functional, like the quantifiers and identity. Nevertheless, these can be construed as functions, and are such

logical vocabulary in each statement which makes the first statement true (more generally, assigns it a designated truth value) makes the second true (designated).

Admittedly, the account of the norms of reasoning I am attempting to sketch is most at home given an at least partly inferentialist account of logical consequence, like that of Peacocke 1987. On such an account, (part of) the meaning of the logical constants will be determined by the inference rules which one must find *primitively compelling* in order to understand the meaning of the logical constant in question.¹¹² On this sort of account, the primitive inference rules for each logical constant constitute (part of) the meaning of said constant, so an agent can be said to be applying the correct rules of deductive reasoning insofar as the agent uses rules which they are primitively entitled to use, or rules that could be justified in terms of the primitive ones, which the agents have some justification for. The second category of rules will thus explain cases where an agent reasons permissibly yet employs rules that are not primitive. For example, mathematicians may, through working in their field, come to derive new proof techniques that are provably valid in terms of previously justified rules. Those mathematicians are then entitled to reason using these rules. We can further allow for cases where an agent becomes justified in applying a derivable rule through a posteriori means: plausibly, agents can become justified in employing certain rules of inference on the grounds of testimony of experts about those rules.

The account is less straightforward on a denotationist approach, since, on such an approach, there is no set of privileged inference rules that characterize logical consequence, and that agents must find primitively compelling. But even if we take the denotationist approach,

given type-theoretic accounts of formal semantics. Note that, on some accounts of logicity, such as that of Tarski 1966, functions generally will count as logical constants (I take no stand on whether this is correct or not).

¹¹² I say, "part of" the meaning because, on Peacocke's account, the primitively compelling inference rules of a logical constant constitute its Fregean Sense. This in turn sets the constant's denotation (given certain other assumptions), but is by itself only part of the meaning of the constant.

there will plausibly be grounds for taking certain sets of rules to be superior or in some sense prior to others. As I discussed in part 1, these will likely be similar to the sorts of considerations that inferentialists use when forming theories of proof-theoretic validity. Even if no set of rules is uniquely determined, the (valid) rules of inference that one *in fact accepts* place restrictions on the rules of inference that they may employ. So even on a wholly denotationalist account, there is hope that we could take certain sets of rules as setting the acceptable rules of deductive inference for an agent, at least in light of the fact that the agent, as a matter of purely contingent fact, accepts certain valid inference rules.

What I've presented thus far is a mere sketch of an account of reasoning. To be fully fleshed-out, we would need to specify a logical theory, including a set of rules that we take to characterize our deductive system. These rules would then set the primitively acceptable rules of deductive reasoning. Ideally, we would also present a theory of non-deductive reasoning. But these are large projects that I cannot hope to pursue in more depth here. Rather, I hope that what I've said thus far sketches out the way that such an account would go.

The Advantages of This View

I've now provided a sketch of a general theory of reasoning. It's time to see what my view has going for it, and why, I believe, we should accept it.

First, my view satisfies the definition of reasoning that I have presented: reasoning is a rule-following activity for transitioning between beliefs in support statements to beliefs in other support statements. This also has the upshot of satisfying Identity*: the rules of reasoning are a special subset of the rules of belief revision. Specifically, they are the rules for forming beliefs in

support statements. Since my account is a rule following account, it construes reasoning as an active, personal level activity. And further, it makes a clear demarcation between deductive and other kinds of reasoning: deductive reasoning involves the formation of support statements whose support relation is deductive, while non-deductive reasoning employs non-deductive support relations.

Second, my account easily handles cases of hypothetical reasoning, and explains what attitude you take toward your conclusions in said reasoning. In hypothetical reasoning, the support statements you infer contain assumptions that you do not necessarily endorse. In supposing said assumptions, you add them to the context of inference, and thus take them for granted. But you at no point need to *believe* said assumptions, nor anything that you derive from them. Rather, what you believe is that the derived conclusions *follow* (in the appropriate sense) from the assumptions and premises that you take for granted. In other words, what you believe are the support statements which you derive on the basis of these assumptions.

Third, my account explains why it is possible to reason correctly and then still fail to form beliefs in the conclusions of your reasoning. When you partake in reasoning of a certain kind, you necessarily form beliefs in support statements of the appropriate kind. In that sense, your belief in the support statement satisfies the norms of belief revision insofar as it satisfies the norms of reasoning. But whether you then revise your attitude in the consequent of a support statement as the norms of belief revision say you ought is another matter, and it is perfectly possible to reason correctly and then fail to adjust your attitude toward the consequent accordingly. The (not necessarily conscious or deliberate) activity of adding and removing beliefs is in general distinct from the ways that said beliefs may be justified, and the considerations on which said beliefs may be based. When you engage in proper reasoning from a

set of premises that you believe, and arrive at a conclusion that you had not previously considered and don't have strong reasons to reject, you typically add that conclusion to your set of beliefs automatically. But it's important to keep in mind that, according to the view I'm espousing, this really is a two-step process: you first reason from P_1, \dots, P_n to C , coming to believe the support statement that $P_1, \dots, P_n \models C$. If you believe P_1, \dots, P_n , you might then automatically add C to your set of beliefs. This is a separate activity from your reasoning, however, and you might have performed the steps of reasoning but failed to come to believe the conclusion.

Fourth, my account can explain the role of reasoning even in cases where the only belief that is at issue is the support statement itself. This will most often happen in mathematical contexts, or in cases of calculation. Consider the following case:

(GRE Math): Terence is taking the GRE. He is asked to solve for x in the equation $4x + 13 = 7 - 2x$. Terence reasons as follows (I add numbering for clarity):

- (1) $4x + 13 = 7 - 2x$ (assumption)
- (2) $6x + 13 = 7$ (adding $2x$ to each side)
- (3) $6x = 7 - 13$ (subtracting 13 from both sides)
- (4) $6x = -6$ (performing $7 - 13$)
- (5) $x = -6/6$ (dividing each side by 6)
- (6) $x = -1$ (performing $-6/6$)

GRE Math is a perfectly standard example of algebraic calculation. But while every step of reasoning that Terence performs is correct, it doesn't make sense to suggest that Terence should

come to believe (unqualified) that $x = -1$. First, there's an issue of making sense of the semantic content of the free variable 'x'. But however that is handled, the more pressing point is that, if Terence encounters another problem which asks him to solve for x, Terence is in no sense justified, despite his present good reasoning, in coming to that new problem with the belief that $x = -1$. My account of reasoning gives a very natural explanation of this: what Terence is really doing is not reasoning about propositions like ' $4x + 13 = 7 - 2x$ ', but rather reasoning about sequents (entailment statements) of the form ' $4x + 13 = 7 - 2x \vDash \dots$ '. That is, Terence assumes that $4x + 13 = 7 - 2x$ given the setup of the problem, and reasons from that to the conclusion $4x + 13 = 7 - 2x \vDash x = -1$. The best representation of Terence's steps of reasoning, then, is as follows:

- (1) $4x + 13 = 7 - 2x \vDash 4x + 13 = 7 - 2x$ (assumption)
- (2) $4x + 13 = 7 - 2x \vDash 6x + 13 = 7$ (adding $2x$ to each side)
- (3) $4x + 13 = 7 - 2x \vDash 6x = 7 - 13$ (subtracting 13 from both sides)
- (4) $4x + 13 = 7 - 2x \vDash 6x = -6$ (performing $7 - 13$)
- (5) $4x + 13 = 7 - 2x \vDash x = -6/6$ (dividing each side by 6)
- (6) $4x + 13 = 7 - 2x \vDash x = -1$ (performing $-6/6$)

Unlike ' $x = -1$ ', Terence can form the unqualified belief in the sequent at step (6) (or any other step), and apply it to further GRE problems he encounters.

Finally, my account vindicates the intuition that the rules of logic are closely related to the rules of reasoning. On my account, for rules of reasoning to be justified for you, they must correspond to rules of deductive or non-deductive logic that you are justified in accepting, whether these are accepted primitively or justified in some other way. Logic *is* normative for

reasoning in a very robust sense (and contra Harman) because it provides constitutive norms for the sort of reasoning in question.

The Norms of Reasoning and Epistemic Norms Revisited

There is one final point that is left to be addressed. As we've seen, my account vindicates Identity*, in that the rules of reasoning double as rules of belief revision for support statements: when one engages in correct reasoning and forms a belief in a support statement, this support statement should be believed provided that the reasoning which was used to arrive at it was good. But there are still some open questions about how my account of reasoning fits into an overall theory of the epistemic. First, nothing I've said has made clear when one should form a belief in the conclusion of an entailment statement that one comes to believe via reasoning, or when one should rather reject a premise. Second, while my account of reasoning makes the rules of reasoning rules of belief revision with respect to support statements, and thus naturally explains what it is to base one's belief in a support statement on one's reasoning, we typically take it that one can base one's belief in non-support statement propositions via reasoning as well. But it isn't clear what this requires on my account.

Regarding the first point, I'll start off by noting that, in order to give a thorough account of when one should reject a belief in some premises or accept a belief in a conclusion as a result of some (good) reasoning, I would need to present an account of how epistemic reasons combine. This is an enormous task, and far beyond the scope of this thesis. But I can nevertheless try to give a sketch of what such an account might look like.

First, it borders on platitudinous that reasoning *transmits justification*. If I believe a set of premises P_1, \dots, P_n , and reason that $P_1, \dots, P_n \models C$, it's plausible that my reasons for believing my premises are transmitted to my conclusion. On my account then, this would have to be reflected in some principle describing how reasons transmit from premises to conclusion provided that I add $P_1, \dots, P_n \models C$ to my beliefs, via reasoning.

But this raises a question: *what are* my reasons to believe a string of individual premises P_1, \dots, P_n ? Plausibly, on any account of epistemic reasons, I will have reasons for or against each individual P_i . But it isn't clear from this exactly what reasons will transfer to C .

Here's a plausible (and non-committal) suggestion: the reasons that transfer to C are the *mereological fusion* of my reasons for each of the P_i . Since the fusion of one's reasons for each of the P_i is plausibly the reason one has for believing the conjunction of P_1, \dots, P_n , it's plausible that what transfers to the conclusion is my justification for the *conjunction* of the premises. This says nothing about the force that my reasons will have on whether or not I should believe C : to establish that, we would need an account of how reasons combine. But there are good reasons to think that reasons are *holistic*: a reason R_1 may, in isolation, be a strong reason to believe P , but in conjunction with a reason R_2 , may provide less support for P than the sum of the supports given by R_1 and R_2 . This holistic character of reasons manifests in cases of undercutting defeat. So plausibly, when I reason from P_1, \dots, P_n to C , the fusion of whatever reasons I have for believing P_1, \dots, P_n transmits to C , and interacts with the reasons I have for or against C . These, in turn, decide whether or not I should believe C .

Second, since reasoning involves forming a belief that $P_1, \dots, P_n \models C$, Steinberger's Directive will plausibly apply when one reasons from the premises to the conclusion. Thus, in such a case, I will have (defeasible) reasons to either believe C or disbelieve at least one of the

P_i. So, upon reasoning from some premises to a conclusion, Steinberger's directive will give me reason to either collectively not believe the premises or believe the conclusion. Which I should do, and which premises I should reject if I should reject any in particular, will depend on our account of how reasons combine, and the reasons I have for each of the premises.¹¹³

So my account of the rules of reasoning seems to sync up nicely with Steinberger's Directive, but even with Steinberger's Directive in hand, we won't have a full epistemic theory unless we have an account of how reasons combine, which I cannot provide here. It's worth saying, however, that adding a support statement to one's beliefs should allow one's reasons to believe the premises to transmit to the conclusion. How strong these reasons are will depend on our theory of how reasons interact.

Turning now to basing, I should note that good reasoning is, on most views, one of many ways to properly base one's beliefs. Beliefs in support statements themselves will be formed directly via reasoning, and will be properly based provided that they are formed through applications of justified rules. But what about my belief in the conclusion of a support statement, were I to form it as a result of reasoning?

Fortunately, it's relatively straightforward to say how reasoning relates to basing on my view. In general, accounts of basing specify what is required for one to take a proposition as a *reason for which* one believes C. To base my belief in C on my reasoning, then, is for my beliefs in the premises P_1, \dots, P_n and the support statement $P_1, \dots, P_n \models C$ to be reasons for which I come to believe C. This basing will be proper insofar as (1) my reasoning which formed my belief in

¹¹³ It's worth briefly reiterating that (as was discussed in part 3), should I reason from no premises to a conclusion, Steinberger's Directive will dictate that I should form the belief in that conclusion. This is as it should be.

the support statement was proper, and (2) insofar as the formation of my belief in C satisfies any other conditions of proper basing that we accept.

This last point highlights a final advantage of my account. As I stated earlier, on most views, it is granted that reasoning *transmits* justification: if you reason from P to Q, and you have strong reasons to believe P, these reasons transfer to Q. But as Balcerak Jackson and Balcerak Jackson (Balcerak Jackson²) point out, there are cases where reasoning seems to *generate* justification.¹¹⁴ Balcerak Jackson² give the following example to demonstrate this:

“Fozzy has no evidence one way or the other concerning the question of whether there are any talking green creatures. But he temporarily adopts the assumption that Kermit is a talking green creature. Working under this assumption, Fozzy thinks to himself: in this case there would be at least one thing—namely, Kermit—who is a talking green creature. Discharging his assumption, Fozzy draws the (material) conditional conclusion that if Kermit is a talking green creature, then there is at least one talking green creature. He goes off in search of Kermit to check his color and language skills”¹¹⁵

Now, Balcerak Jackson²'s example is problematic: if it's true that Fozzy has no evidence one way or another about talking green creatures, then what exactly does Fozzy take the referent of the name 'Kermit' to be? But we can generate a similar example without postulating that our agent supposes a named entity whose existence is in question. Consider the following, rather banal case:

John is wondering whether his friend Andrew is at the store. Playing around with logic, he temporarily adopts the assumption that Andrew is at the store. Working under this assumption, John thinks to himself: in this case there would be at least one person – namely Andrew – who is

¹¹⁴ Balcerak Jackson and Balcerak Jackson 2013.

¹¹⁵ Balcerak Jackson and Balcerak Jackson 2013, p. 116.

at the store. Discharging his assumption, John draws the (material) conditional conclusion that if Andrew is at the store, then there is at least one person at the store.

John's newfound belief that if Andrew is at the store then someone is, while not incredibly interesting, is perfectly justified. How can this be, if it proceeds from no premises? On my account, John's reasoning is as follows:

- (1) Andrew is at the store \models Andrew is at the store
- (2) Andrew is at the store \models At least one person is at the store (from 1 by $\exists I$)
- (3) $\emptyset \models$ If Andrew is at the store, then at least one person is at the store. (from 2, by $\rightarrow I$)

John's belief in each of these support statements is justified precisely because each is formed via application of a rule of reasoning that John is justified in using. By Steinberger's Directive, (3) entails that John has reason to see to it that, if he believes each of the (empty set of) P_i , then he believes 'If Andrew is at the store, then at least one person is at the store'.¹¹⁶ Since the antecedent is vacuously true, John has (decisive) reasons to believe this conclusion. If John comes to believe said conclusion, and further bases his newfound belief on the support statement (3), then he plausibly will be doxastically justified in believing 'If Andrew is at the store, then at least one person is at the store'.

So if it's granted that one may be justified in forming beliefs in support statements via application of one's justified rules of reasoning, then my account, in conjunction with Steinberger's Directive, demystifies how reasoning from no premises can generate justification in conclusions. It also allows us to make sense of how one could base one's belief on one's

¹¹⁶ I assume that, in having formed a belief in (3), John satisfies the condition "and one considers C" in SD, or the condition, "and one is aware of the question of whether C" in SD*.

reasoning: to do so is to take a support statement (formulated via reasoning) *as a reason* to believe the conclusion.

Conclusion to Part 4

In this section, I've discussed the nature of reasoning, and presented a novel account of reasoning, drawing from the ideas of Boghossian and Tichý, among others. I then argued that this account has many advantages, and satisfies numerous desiderata that a theory of reasoning should satisfy. I then discussed how my account of reasoning should fit in to a full theory of the epistemic. While the account I've presented is a mere sketch, and sidesteps many difficult issues, I hope to have made it seem plausible that my account could be turned into a plausible and comprehensive theory of reasoning, which interfaces nicely with our other epistemic norms.

Most importantly for the topic of this thesis, I've argued that (deductive) logic is normative for reasoning in that it generates constitutive norms that (deductive) reasoning should satisfy. In particular, I've argued that the rules of natural deduction that one is justified determine the rules of deductive reasoning which are acceptable: precisely those rules which correspond to the rules of natural deduction. My hope is that a similar story could be told linking the rules of non-monotonic logic with the rules of defeasible reasoning. But if my account is correct, then logic is normative for reasoning in a very robust sense. Harman and other opponents of the normativity of logic have failed to see this, because they have been operating under a mistaken conception of what the subject matter of reasoning is. It's certainly true that reasoning from P to Q does not straightforwardly require one to form a belief in Q. But the reason for this is that your reasoning only secures belief in the support statement that $P \models Q$. Deriving this support

statement through any correct rule of logic *does* straightforwardly license you to come to believe it. But whether you should also come to believe Q, or rather reject P, will depend on your reasons for each. For propositions of this sort, the rules of belief revision are distinct from the rules of reasoning.

Part 5: Integrating the Normativity of Logic into a Theory of Epistemic Normativity

In this section, I discuss how the epistemic norms generated by logic that I have defended in Part 3 and Part 4 fit into an overall epistemological theory. In particular, I investigate how said norms interact with Evidentialism, which I introduced in Part 2. I argue that reasons and oughts generated by Evidentialism are best construed as directives. I then investigate the relationship between Evidentialism and Steinberger's Directive.

Recap: The Norms of Logic

Before investigating how the norms of logic interact with other epistemic norms, it will be worth revisiting the set of norms that have been defended in Part 3 and Part 4. In Part 3, I investigated what norms, if any, logic generates on *belief*. Drawing heavily from work by Florian Steinberger, I argued that there were three kinds of norms which logic might generate: evaluations, directives, and appraisals. Using Steinberger's analysis and the desiderata that Steinberger, MacFarlane, and other authors have considered, I argued that logic generates at least the following two norms:

(WO+) If $A_1, \dots, A_n \models C$, then S ought to (believe C, if S believes A_1, \dots, A_n)

(SD) "If according to S's best estimation at the time, S takes it to be the case that $A_1, \dots, A_n \models C$ and S has reasons to consider or considers C, then S has reasons to (believe C, if S believes all of the A_i)"

WO+ is a wide-scope closure requirement, which requires of agents that they have beliefs which are closed under logical consequence. As we saw in Part 3, this principle is the most plausible

evaluation that logic could generate, as it satisfies all of the desiderata that an evaluation should satisfy that we have thus far considered. Similarly, SD, or Steinberger's Directive, is the bridge-principle that best satisfied all of the desiderata that a plausible directive should satisfy. It's important to recall that the strength of the *reasons* operator in SD will vary given the context in order to account for the Preface Paradox: plausibly, the more premises in A_1, \dots, A_n , the weaker the reasons generated by Steinberger's Directive. This allows said reasons to be overridden in preface-like cases, where the set of premises is large and C is the conjunction of said premises.

I also considered a modified version of SD, which was meant to ensure that Steinberger's Directive would be compatible with epistemic views that rejected that there could be epistemic reasons to consider a proposition:

(SD*) If according to S's best estimation at the time, S takes it to be the case that $A_1, \dots, A_n \models C$ and S is aware of the question of whether C, then S has reasons to (believe C, if S believes all of the A_i).

So WO+ describes a plausible logic-generated evaluation, and SD or SD* provides a plausible logic-generated directive. As for appraisals, there will be many possible principles linking the facts of logic to facts of epistemic blameworthiness and praiseworthiness, which vary in strength between that of SD (or SD*) and that of WO+, depending on their degree of idealization. Thus, Steinberger gives us a plausible set of epistemic norms on belief which are generated by logic, and which satisfy every desideratum we've considered.

In Part 4, I discussed the norms of logic on *reasoning*. Drawing from work by Paul Boghossian and Pavel Tichý, among others, I sketched an account of reasoning according to which reasoning is a rule-following endeavor (which thereby satisfies the *taking-condition*), and

that the rules of reasoning are rules for forming beliefs in *support statements*: propositions which say that a set of premises either deductively or defeasibly entail a conclusion. I showed the virtues of my account, and argued that (deductive) logic would generate constitutive norms on (deductive) reasoning. These norms would be satisfied provided that the agent – when reasoning deductively – reasons using the natural deduction rules of logic that the agent is justified in accepting (either primitively or otherwise). On my account then, the rules of logic are normative for the rules of reasoning in a very robust sense, but the rules of reasoning do not by themselves determine the rules of belief revision (except when the content of the belief being formed is itself a support statement).

What Kind of Norm is Evidentialism?

On my account then, the norms discussed in the previous section are the (epistemic) norms of logic. But now it's worth asking: how do the norms of logic relate to other epistemic norms? In particular, suppose we adopt *Evidentialism* as our overarching theory of what one ought to believe:

Evidentialism about the Epistemic Ought: Whether S ought to believe P is determined by whether S has, at the time, adequate evidence for P. Assuming S is aware of the question of whether P, then (1) if S has adequate evidence for P, then S ought to believe P, and (2) if S does not have adequate evidence for P, then S ought to not believe P (i.e., suspend on P or disbelieve P).

This definition of Evidentialism, which was given in Part 2, states a synchronic norm on *belief*. It is a result of combining Evidentialism about Epistemic Reasons with the Uniqueness thesis,

which stated that for any body of evidence, only one doxastic attitude will be permissible for each proposition. An adherent of Evidentialism will further accept the following definitions of propositional and doxastic justification:

Propositional Justification: An agent S has (propositional) justification for a proposition P at time t iff S's (*total*) *evidence* at t supports belief in P.

Doxastic Justification: an agent S is (doxastically) justified in a belief in a proposition P iff:

(7) S has propositional justification for P

(8) S believes P

(9) S *properly bases* their belief in P on their *evidence* for P.

At this point, it's worth asking exactly what kind of norm the Evidentialist 'ought' purports to be. I claim that Evidentialism, as stated, is best viewed as a directive, rather than an evaluation or appraisal (though it may be accompanied by an appropriate set of appraisals for blameworthiness and praiseworthiness). There are problems with positing Evidentialism as an evaluation: that is, as an ultimate standard of correctness of belief. One problem is that Evidentialism as such does no justice to veritistic intuitions. There is no shortage of philosophers who have argued that "truth is the aim of belief", or qualified versions thereof, and Evidentialism cannot seem to accommodate these intuitions if evidence-responsiveness is the ultimate standard of evaluation for belief. The reason for this is that, on almost all conceptions of evidence, it is possible for a proposition to be all-things-considered supported by the evidence an agent has, while still being false. But if having beliefs which were supported by the evidence was the *ultimate* standard of

correctness of belief, then there would be no sense in which one's beliefs would be better or more correct provided that they were true or known, as opposed to merely justified.¹¹⁷

Now, it's important to note that the argument I just gave for why Evidentialism should be construed as a directive does not preclude the possibility of giving a veritistic definition of evidence, or of evidential value.¹¹⁸ Plausibly, in any domain, the true directives are those norms which are effective for giving guidance toward whatever actions (or beliefs) the true evaluations recommend. If I ought – in the *evaluative* sense – do whichever of my options will maximize my own happiness, then I plausibly ought – in the *directive* sense – do whichever option seems to have the greatest expected happiness. Plausibly, that will involve taking an action which is typically happiness-conducive (to the greatest degree). Similarly, if the true epistemic *evaluations* recommend having only true or known beliefs, then the true *directives* will be norms which recommend believing in ways that are typically conducive to my beliefs being true/known. Plausibly, this means believing what is best supported by my evidence. So there is a strong case to be made that the true directives will recommend I believe according to my evidence. We can thus give a veritistic vindication of Evidentialism *as a directive*. But this will not thereby give a vindication of Evidentialism *as an evaluation*, since believing according to the evidence, while conducive to having epistemically valuable (i.e. true or known) beliefs, does not

¹¹⁷ Similar arguments can be made to demonstrate that truth cannot be the ultimate evaluative standard of belief, because plausibly, knowledge is superior to mere true belief. While I'm sympathetic to such a view, said view hinges on there being *sui generis* epistemic value (alternatively, objective epistemic reasons) to have justified beliefs, so I won't commit to such a view here. More importantly, the argument is much more acute when applied to views which take justified belief, but not truth, to set the ultimate evaluative standard, because such views preclude defending the value of justification in terms of truth-conduciveness. The argument I give in this paragraph is a variant of the argument known in the literature as *the Swamping Problem*. See Stapleford 2016 for a discussion of the problem and attempted rebuttal to the kind of argument I make here.

¹¹⁸ Though we might not necessarily want to do that either. See Berker 2013a and Berker 2013b for arguments that the question "what should I believe?" should not be answered in terms of truth-conduciveness. I take no stand on whether these arguments are correct here.

guarantee that my beliefs will be as epistemically valuable as they could be (i.e., it will not guarantee that all my beliefs are true/known).

If I'm correct, then Evidentialism should be construed as a theory about epistemic directives, not epistemic evaluations. Of course, one might object to my argument in numerous ways.¹¹⁹ I'll briefly consider a couple of objections. First, one might object to my claim that Evidentialism should be construed as a directive by asking the following question: Why can evidentialism be a good directive if it's a poor evaluative norm? If knowledge is the best evaluative norm, why should agents not be directed to get knowledge?

This question misunderstands the nature of directives. The point of a directive is to *give guidance* to an agent. Even if 'believe only what you are in a position to know' is the correct epistemic evaluation, this norm can't serve as a successful directive, simply because it fails to be guidance giving in the same way that, 'believe what is true' fails to be guidance giving. Plausibly, the closest you can come to knowing whether you are in a position to know a proposition is to know that you have decisive evidence in favor of that proposition. But this just is the directive interpretation of Evidentialism.

Second, one might object that the directive interpretation of Evidentialism is untenable, because what one's evidence supports extends beyond what a directive could require one to believe. If, for example, it is too difficult for an agent to see their evidence supports P, then Evidentialism will nevertheless require them to believe P, even though no true directive could demand that the agent believe P.

¹¹⁹ Thanks to Julia Staffel for pointing out these objections to me.

In order for this objection to have force, we would need a markedly externalist conception of evidence: it would have to be possible for an agent to have evidence for P, but be unable to recognize that said evidence justifies belief in P. I don't think such externalist conceptions of evidence are plausible, in part because I think evidence must be construed hyperintensionally (for reasons I'll discuss later), and partly because of general considerations in favor of epistemic internalism. But to give a full-fledged defense of internalism is far outside of the scope of this thesis. Anyone who finds this objection compelling can thus take the following discussion as investigating the following question: if evidence *were* internalistic, then how would the norms of logic relate to Evidentialism?

Evidentialism and Steinberger's Directive

If I'm correct, then we should construe Evidentialism as a directive. If so, then Evidentialism is in the same class of epistemic norms as Steinberger's Directive. So we might ask: what the relationship is between the two? Is one reducible to the other? And is it possible to defend both simultaneously?

Recall that Evidentialism, as stated above, was derived in Part 2 from two theses:

Evidentialism about Reasons: Epistemic reasons for believing propositions are constituted by an agent's *evidence* for/against those propositions.

Uniqueness: For any set of epistemic reasons, there is exactly one attitude which is permissible toward each proposition. Hence, that attitude is obligatory.

It's worth noting that Evidentialism about Reasons is actually a fairly strong normative claim, in that it *identifies* evidence and epistemic reasons. But one might think that, while evidence is a

perfectly good sort of epistemic reason, there could be others. Perhaps the class of epistemic reasons is strictly broader than the class of things that count as evidence. Call this view *Weak Evidentialism*:

Weak Evidentialism: (*At least*) some of an agent's epistemic reasons for believing propositions are constituted by an agent's *evidence* for/against those propositions.

Weak Evidentialism is eminently plausible. And if Weak Evidentialism is true, then there will be (evidential) epistemic norms governing how one is permitted to/ought to believe in light of one's evidence.¹²⁰ So with this in mind, we might ask a general question: How do the norms of logic (in particular, Steinberger's Directive) relate to the true evidential norms? There are a few possibilities, which mirror the possibilities for accounting for the relationship between evidential norms and coherence norms discussed in Worsnip 2018.¹²¹ There are four options:

- (1) The norms of logic are reducible to evidential norms
- (2) The norms of evidence are reducible to the norms of logic
- (3) Neither the norms of logic nor the norms of evidence are reducible to each other, and there are true norms of both kinds.

¹²⁰ Specifying what norms these are is a large project, comparable to the project of specifying how reasons combine, and so I won't undertake this project here. But as a mere example, a true evidential norm might be: "If you ought believe P, and your evidence for P is just as strong as your evidence for Q, then you ought believe Q."

¹²¹ There is a good reason that the possibilities for the relationship between the norms of logic and the norms of evidence mirror the possibilities for the relationship between coherence norms and evidential norms: plausibly, the norms of logic are a kind of coherence norm. Coherence norms are typically construed as norms which restrict which attitudes may be rationally held simultaneously.

(4) Neither the norms of logic nor the norms of evidence are reducible to each other, but there are only norms of one kind.¹²²

However, while these four options are all possible in theory, it's easy to see that option (2) is false. Option (2) amounts to the claim that Steinberger's Directive entails the true norms of evidence. But we've already seen that this cannot be the case because Steinberger's Directive is (by design) wide-scope: when, according to one's best estimation, $A_1, \dots, A_n \models C$, and one is aware of the question whether C /considers C , Steinberger's Directive gives one reasons to either not believe all of the premises or believe the conclusion. But Steinberger's Directive gives no verdict on which one should do. So we will need other epistemic norms which can make this decision, and plausibly, these will be evidential norms. This also means that, if option (4) is true, it is true because there are *only* evidential norms on belief, rather than because there are only logical norms on belief.

In light of the previous discussion and the eminent plausibility of Weak Evidentialism, we can identify the following live options:

(Reducibility) The norms of logic are reducible to evidential norms

(Non-Reducibility+) Neither the norms of logic nor the norms of evidence are reducible to each other, and there are true norms of both kinds.

(Non-Reducibility-) Neither the norms of logic nor the norms of evidence are reducible to each other, but there are no true logical norms.

¹²² For this discussion, I set aside the possibility (which to my knowledge, remains undefended) that there are no true evidential norms and no true norms of logic.

With these three options in mind, we can return to the question of how Evidentialism and Steinberger's Directive relate. Reducibility and Non-Reducibility- (henceforth NR-) are compatible with Evidentialism, while Non-Reducibility+ (henceforth NR+) is incompatible with Evidentialism. The reason for this is that NR+ entails that there are epistemic reasons which are not evidence. Hence, Evidentialism about Reasons will be false, and by extension, so will Evidentialism about the Epistemic Ought. But even so, Weak Evidentialism will still be true.

There's a caveat to the conclusion of the last paragraph. Even if NR+ is true, it may still be that Evidentialism about the Epistemic Ought is *extensionally adequate*, despite not being true. The reason is the following: Evidentialism about the Epistemic Ought says that what one ought to believe is determined (solely) by one's evidence. If NR+ is true, then this claim is false: the norms of logic generate non-evidential reasons for belief. But it still may be the case that, in every case, what one's evidence recommends and what one ought to believe coincide. In such a scenario, non-evidential reasons exist, and contribute to what one ought to believe, but they never change the verdict that one's evidence recommends: what one all-things-considered ought to believe is just what one evidentially ought to believe, in every case. If this were correct, then, while NR+ would be true (so we've supposed), the norms of logic would be in some sense idle.

The truth of Evidentialism given Reducibility or NR- hinges on whether there are any other kinds of epistemic norms, other than logical and evidential norms. For example, it is possible that there are *probabilistic coherence norms*. If so, then Evidentialism will be true provided that these are reducible to evidential norms. But if there are no other plausible epistemic norms then given either NR- or Reducibility, Evidentialism will be true.

With all of this in mind, I'll now turn to the arguments for and against each of these options, starting with Reducibility.

Is Steinberger's Directive Reducible to Evidentialism?

In light of the discussion in Part 3, we found that the most plausible candidate for a logical directive for belief was Steinberger's Directive (construed either as SD or SD*). For Reducibility to hold, it would have to be the case that it is possible to derive Steinberger's Directive from Evidentialism. In order for that to be the case, the epistemic reasons imputed by Steinberger's Directive would have to be *evidence*.

But it's very hard to see how this could be so. First, it's a truism about evidence that evidence either supports or goes against individual propositions. If, for example, we construe evidence probabilistically, a body of evidence T will be evidence for P iff T *raises* the probability of P relative to P's background probability, and T will be evidence against P iff T *lowers* the probability of P relative to P's background probability.¹²³ But the epistemic reasons imputed by Steinberger's Directive do not do this. Given Steinberger's Directive, if $A_1, \dots, A_n \models C$ holds in an agent's best estimation and the agent considers C (or is aware of the question of whether C), then the agent has reasons to not both believe all of the A_i and not believe C. But these reasons do *not* point in favor of C or against C, nor do they point in favor of/against any or all of the A_i : they are entirely silent on which option the agent should choose. There is good reason for this: as was discussed in Part 3, if the reasons *did* point in a particular direction, then Steinberger's Directive would not be flexible enough to accommodate situations where the agent should not believe all of the premises, as well as situations where the agent should come to believe the conclusion.¹²⁴ And so it's hard to see how the reasons imputed by Steinberger's

¹²³ What sets this background probability is open for debate. On a subjective Bayesian construal, P's background probability for an agent will be the *prior* probability that the agent assigns to P before conditionalizing on E.

¹²⁴ The preceding argument was discussed in Part 3 as a desideratum on true bridge principles, and was named *Belief Revision*. The argument originates from Harman 1986.

Directive could be evidence, since it isn't clear what proposition they'd be evidence for. It isn't the conclusion C, but nor is it any of the premises A_i .

If the discussion in the preceding paragraph is correct, then the prospects for *Reducibility* seem slim. While no argument has been given that decisively rules out the possibility that the reasons imputed by Steinberger's Directive could be evidence, it's hard to see how this could be the case, because it's hard to find a plausible candidate for what the reasons could be evidence for. This means that the norms of logic are likely not reducible to evidential norms. We are thus left with NR+ and NR- as candidates for how Steinberger's Directive relates to Evidentialism. That is to say, either Evidentialism is true and Steinberger's Directive is false (NR-), or Steinberger's Directive is true, as is Weak Evidentialism, but Evidentialism is false (NR+).

The Case For NR-

A strong argument in favor of NR- (and thus against NR+) can be reconstructed by drawing from some work by Niko Kolodny.¹²⁵ According to Kolodny, from a deliberative standpoint (i.e., with respect to directives), there is no room for wide-scope norms (coherence norms) such as Steinberger's Directive. Kolodny makes the following conjecture:

[T]here are no requirements of formal coherence as such... Requirements of formal coherence seem plausible, I suggest, because they would explain certain 'violation claims'—to the effect that when our attitudes are formally incoherent we violate a requirement—and certain 'satisfaction claims'—to the effect that when we make our attitudes formally coherent we satisfy a requirement. We may be able to explain several of the violation claims, I suggest, by appealing instead to 'what reason requires': that is, to what we ought to believe or choose given what matters

¹²⁵ Kolodny 2007.

from the standpoint of deliberation. The attitudes that reason requires, in any given situation, are formally coherent. Thus, if one has formally incoherent attitudes, it follows that one must be violating some requirement of reason. The problem is not, as the idea of requirements of formal coherence as such suggests, that incoherent attitudes are at odds with each other. It is instead that when attitudes are incoherent, it follows that one of these attitudes is at odds with the reason for it—as it would be even if it were not part of an incoherent set.¹²⁶

Kolodny considers numerous examples of plausible seeming coherence directives, argues against each, and presents an error theory for why we would be inclined to think that each holds. Among the putative coherence directives that Kolodny considers are the following:

Non-Contradiction (N): One is rationally required (if at *t* one believes *P*, then at *t* one does not believe not-*P*).¹²⁷

Single-Premiss Closure (C): When *Q* is a logical consequence of *P*, one is rationally required (if at *t* one believes *P*, then at *t* to believe *Q*).¹²⁸

According to Kolodny, we are inclined to think that these are true directives because we hold the following intuitions:

Violation Claim about Non-Contradiction (VN): If one believes at *t* that *P* and believes at *t* that not-*P*, then one violates some requirement.¹²⁹

Violation Claim about Single-Premiss Closure (VC): When *Q* is a logical consequence of *P*, if at *t* one believes *P*, but does not believe *Q*, then one violates some norm.¹³⁰

As well as the further intuition:

¹²⁶ Kolodny 2007, p. 231.

¹²⁷ Kolodny 2007, p. 229.

¹²⁸ Kolodny 2007, p. 249.

¹²⁹ Kolodny 2007, p. 232.

¹³⁰ Kolodny 2007, p. 249.

Violation Intuition about Logicality (VL): If not- P_n is a logical consequence of P_1 and P_2, \dots, P_{n-1} , then if at t one believes P_1 , believes P_2, \dots , and believes P_n , then one violates some requirement; and if P_n is a logical consequence of P_1, P_2, \dots, P_{n-1} , and at t one believes P_1 , believes P_2, \dots , and believes P_{n-1} , but does not believe P_n , then one violates some requirement.¹³¹

(N) and (C) explain (VN) and (VC), respectively. However, Kolodny argues that (VN) and (VC) are also adequately explained by the following putative facts about how epistemic reasons (evidence) combine, respectively:

First Comment on Reason Patterns (R1): In any given case, either one lacks sufficient reason to (i.e., reason requires one not to) believe P , or one lacks sufficient reason to believe not- P .¹³²

Second Comment about Reason Patterns (R2): When Q is a logical consequence of P , either one lacks sufficient reason to believe P , or one has conclusive reason to believe Q .¹³³

(R1) explains (VN), because if one believes both P and $\sim P$, then, given (R1), one is believing beyond one's evidence: either one lacks sufficient evidence for P , or one lacks sufficient evidence for $\sim P$. Hence, (VN) is true: believing P and $\sim P$ simultaneously does violate a requirement. (R2), in turn, explains (VC). The reason is that, if $P \models Q$, then given (R2), one either lacks sufficient evidence to believe P , or one has such evidence, and hence has sufficient evidence to believe Q . So one violates a requirement by believing P and not believing Q , but this is an *evidential* requirement, not a *coherence* requirement. As for (VL), Kolodny claims that, in light of the Preface Paradox, our intuitions about (VL) should simply be resisted: generally speaking, there need be nothing wrong with believing a set of premises (as the evidence

¹³¹ Kolodny 2007, p. 253.

¹³² Kolodny 2007, p. 235.

¹³³ Kolodny 2007, p. 249.

demands), but failing to believe something that follows from them jointly (say, their conjunction), as the Preface Paradox shows.

Kolodny goes on to give a general error theory that purports to explain all of the violation-intuitions that seem to favor coherence norms (or at least, all of the violation intuitions that Kolodny thinks we should keep), and concludes that there are no true coherence norms. If Kolodny is correct, this bodes poorly for Steinberger's Directive, and points in the direction of NR-. There are two reasons for this. First, Steinberger's Directive is itself a coherence norm, which Kolodny would likely purport to be able to explain. Second, given certain weak assumptions about the structural properties of evidence, Steinberger's Directive will entail (appropriately qualified versions of) both (N) and (C), as I'll now show:

First, regarding (C), recall that (C) states that, if $P \models Q$, one ought to (believe Q, if one believes P). This is already quite similar to Steinberger's Directive, as Steinberger's Directive states that, if $P \models Q$, and one considers/has reasons to consider (or is aware of the question of whether) Q, then one has (contextually strong) reasons to (believe Q, if one believes P). But since P is an almost maximally un-prefacelike set of premises, these reasons will plausibly have the force of an ought. Thus, Steinberger's Directive vindicates an attitudinally restricted version of (C), which is significantly more plausible as a directive anyway: even if there are true coherence directives, one could hardly be faulted for not having beliefs which are closed under single-premise closure, for the same reason that one could hardly be faulted for failing to have beliefs which are closed under logical consequence.¹³⁴

¹³⁴ There's a simple procedure for translating any instance of logical consequence based on any premises into a case of single-premise entailment. If $P_1 \cdots P_n \models C$, then the conjunction of each of the P_i , $P_1 \& \cdots \& P_n$, is a single premise which also entails C.

(N) follows from Steinberger's Directive given certain weak structural constraints on how evidence combines. (N) states that an agent S ought to (not believe $\sim P$, if S believes P).

Given Steinberger's Directive, the following holds:

(Contra) If according to S's best estimation at the time, S takes it to be the case that $P, \sim P \models P \& \sim P$ and S has reasons to consider or considers $P \& \sim P$, then S has reasons to (believe $P \& \sim P$, if S believes both P and $\sim P$)

Suppose the antecedent is satisfied. Then S has reason to (believe $P \& \sim P$, if S believes both P and $\sim P$). But again, the set of premises are highly un-prefacelike, so these reasons plausibly have the force of an ought. Hence, S ought to (believe $P \& \sim P$, if S believes both P and $\sim P$). But it is a plausible structural constraint on how evidence combines that one can never have decisive evidence for a blatant contradiction like $P \& \sim P$.¹³⁵ That is, it is a plausible constraint on evidence that S ought not believe $P \& \sim P$. Thus, the only way to satisfy Steinberger's Directive while believing according to one's evidence is to not both believe P and believe $\sim P$. So S plausibly ought not (believe P and believe $\sim P$). Taking 'if...then' to be the material conditional (as we have been), it then follows that S ought (not believe $\sim P$, if S believes P). Hence, in any case in which the antecedent of Steinberger's Directive is satisfied, a version of (N) holds. And if it's further the case that the entailment fact $P, \sim P \models P \& \sim P$ could not fail to hold in an agent's best estimation, then (N) will hold whenever an agent considers both a proposition and its negation.

¹³⁵ Note that Kolodny could not plausibly contest this, as it plausibly follows from R1. This assumption is also encoded in many mainstream theories of evidence, particularly probabilistic theories, according to which, the probability of any contradiction must be 0. Even if this constraint is too strong, I think most will agree that *blatantly* contradictory propositions like $A \wedge \sim A$ will never have decisive evidence in their favor. Dialethists of course might disagree with this, but a Dialethist will not accept (N) as a true coherence norm either.

So we've derived qualified versions of (N) and (C) from Steinberger's Directive. If there are good reasons to think that neither of these norms hold, and that (R1) and (R2) are true principles governing evidence which satisfactorily explain the intuitions which motivate (N) and (C) (namely, (VN) and (VC)), then this would be a strong reason to reject Steinberger's Directive in favor of Evidentialism. To evaluate whether this is in fact the case, we should evaluate the cases for (R1) and (R2).

(R1) states that, in every case, either one ought to not believe P, or one ought not believe \sim P. That is, in every case, either one lacks sufficient evidence for P, or one lacks sufficient evidence for \sim P. This seems highly plausible, because it follows from the following principle:

Stronger Evidence (SE): One has sufficient reason (i.e., reason permits one) to believe P only if the evidence indicates that P more strongly than it indicates that not-P.¹³⁶

Suppose an agent has sufficient evidence for P. Then by (SE), said agent will have stronger evidence for P than \sim P. If we grant that for an agent to have sufficient evidence for P, said agent must have more evidence for P than against P, and we further grant that evidence for \sim P is evidence against P, then an agent with sufficient evidence for P will lack sufficient evidence for \sim P. Thus, the second disjunct of our statement of (R1) above will be true. And if one lacks sufficient evidence for P, then the first disjunct of our statement of (R1) is true. So (R1) follows from eminently plausible principles about evidence.

(R2) is more problematic. (R2) states that, if $P \neq Q$, one either lacks sufficient evidence for P or has sufficient evidence for Q. Kolodny derives (R2) from Uniqueness and a principle Kolodny calls (ET):

¹³⁶ Kolodny 2007, p. 253. This principle is rather plausible, and is accepted by most Evidentialists.

Evidence Transmission (ET): When Q is a logical consequence of P, the evidence that Q is at least as strong as the evidence that P.

So according to (ET), if $P \models Q$, then one's evidence for P must be at least as strong as one's evidence for Q.

A Brief Aside: The Hyperintensionality of Evidence

(ET) is a popular thesis: it is implicit in most mainstream theories of evidence, particularly in probabilistic theories (such as any form of Bayesianism). Nevertheless, I think (ET) is false, because I think there is a very strong case to be made that evidence is hyperintensional. That is, it is possible to have evidence for P without having evidence for Q, despite the fact that $P \models Q$. While giving a full defense of the hyperintensionality of evidence is beyond the scope of this thesis, it will be worth giving a truncated defense of it to at least clarify why anyone would hold it.

First, there are linguistic reasons to think that claims about evidence exhibit substitution failure. Consider the following arguments:

(Marilyn Monroe) Max has evidence that Marilyn Monroe is Marilyn Monroe.

Marilyn Monroe is Norma Jean Baker. Therefore, Max has evidence that Marilyn Monroe is Norma Jean Baker.¹³⁷

¹³⁷ This example will not work given certain (particularly Neo-Russellian) theories of hyperintensionality, because on those theories, the hyperintensionality of propositions stems from the structure of the propositions. But 'Marilyn Monroe' and 'Norma Jean Baker' are both names, whose only structural constituent is the object they denote. However, even for those theories, a counterexample that will work can be constructed by substituting 'the individual who is identical to Marilyn Monroe iff no one loves all and only those who do not love themselves' for 'Norma Jean Baker'. This complex description necessarily picks out Marilyn Monroe, but because it contains structural components

(Math Substitution) John has evidence that 3 is the positive solution to $x^2 - 9 = 0$.
3 is the positive solution to $x^2 + 4^2 = 5^2$. Therefore, John has evidence that the
positive solution to $x^2 + 4^2 = 5^2$ is the positive solution to $x^2 - 9 = 0$.

There is a strong intuition that both of these arguments are invalid. Intuitively, Max does not necessarily have evidence that Marilyn Monroe is Norma Jean Baker merely because he has evidence of Marilyn Monroe's self-identity. Similarly, solving the equation $x^2 - 9 = 0$ need not give you evidence that the solution to said equation is also the solution to a different equation, namely $x^2 + 4^2 = 5^2$. If this is right, then evidence is hyperintensional, because substituting intensionally equivalent expressions into sentences of the form "S has evidence that P" introduces substitution failure. So there is good linguistic evidence for the hyperintensionality of evidence.

Second, there are very counterintuitive consequences for construing evidence intensionally (or extensionally) rather than hyperintensionally. The worst of these, in my view, is the following. If evidence is not construed hyperintensionally, then no proposition or set of propositions will ever count as raising the evidential support for any logical truth, since every agent will always have maximal evidence for any logical truth, no matter what propositions they believe or evidence they have. That is to say, given intensionalism about evidence, *learning the entire set of propositions which constitutes the proof of a logical truth doesn't gain you any evidential support for the logical truth whatsoever*. If we further define a piece of evidence for a proposition P to be any proposition which, upon learning it, raises the amount of evidential support one has for P, we will have to say that nothing counts as evidence for logical truths, *even*

that 'Marilyn Monroe' does not (e.g. iff, the loving relation, etc), it will have a different hyperintension on any theory of hyperintensionality.

their proofs.¹³⁸ This seems patently absurd: if anything counts as evidence for a proposition, surely a *proof* of that proposition should count.

Returning to the Case for NR-

The preceding discussion, while brief, should hopefully make it plausible that evidence should be construed hyperintensionally. If so, then (ET) is false: you are not guaranteed to have at least as strong evidence for any consequence Q of P as you have for P . This then casts doubt on (R2), as it is possible to have sufficient evidence for P and lack sufficient evidence for Q despite $P \models Q$.

Unfortunately, this observation does not undermine Kolodny's argument that (R2) satisfactorily explains our intuition (VC). The reason is that (VC), in its current form, is not plausible. (VC) states that an agent violates some norm if they believe P but fail to believe Q when $P \models Q$. But as a *directive*, this is plausibly false: you only violate some norm if you further *recognize* the entailment $P \models Q$. But there is a principle in the same vein as (R2) which explains this intuition: if $P \models Q$ and S recognizes this, then S either lacks sufficient evidence for P or has sufficient evidence for Q . This principle is plausibly true, and satisfactorily explains the qualified version of (VC).

So (at least) qualified versions of (R1) and (R2) seem eminently plausible. These would explain our violation-intuitions (VN) and (VC), or at least the suitably qualified versions of these

¹³⁸ Note that, if we manage to avoid this by giving a definition of *being a piece of evidence for P* that does not require raising P 's probability, then we will have to say that *every* proposition is evidence for P , when P is a logical truth. This is because any proposition whatsoever, including, say, the proposition that you drank coffee this morning, entails P if P is a logical truth.

intuition claims. This leaves room for Kolodny's error theory against coherence directives. But it doesn't in any way force us to believe that (N) and (C) are false. Why not simply think that (N) and (C) hold as well?

One reason is methodological. We shouldn't postulate norms that don't do any work. If suitably qualified versions of (R1) and (R2) fully explain our intuitions that support (N) and (C) (or the suitably qualified versions thereof), then it would appear that there is no reason to postulate coherence norms in addition.

But there is another, more important reason to be skeptical that (N) (or the qualified version of it that follows from Steinberger's Directive) is a true directive. (N) differs from (R1) in the following respect: one can satisfy (N) (but not (R1)) by believing *against* one's evidence. (N) merely requires you to not simultaneously believe P and \sim P. If your evidence supports P, and hence doesn't support \sim P, you still satisfy (N) if you believe \sim P. R1, by contrast, will not be satisfied in this case, since according to R1, you ought to believe P.

Notice that, as it stands, this observation seems to *support* the existence of coherence norms like Steinberger's Directive and (N). We tend to think that agents who believe against their evidence but who are consistent in doing so are in some sense better off than those that aren't. While we wouldn't want to say that such an agent was believing *rationally*, we might say something like, "well, at least they're consistent" when appraising such an agent. Having consistent beliefs which are counter to the evidence seems in some sense exculpatory of blame, even if it doesn't render one's beliefs rational. So if there is a close tie between appraisals and directives, this might indicate that there are in fact true coherence directives, like (N).

So the problem is not that (N) imputes reasons that (R1) does not. The problem is rather that it's very hard to see how these reasons could justifiably play a role in your deliberation, because it's hard to see what could be said from a deliberative standpoint for revising your beliefs against the evidence. As Kolodny puts it, "If nothing can be said for satisfying P even in this way—if something can be said only for satisfying N with reason—then R1 would appear to be the sole normative truth in the vicinity."¹³⁹ Kolodny goes on to address a number of potential grounds for accepting (N), and finds them all wanting. We are thus left without an explanation of why (N) makes for a plausible directive, over and above (R1).

There is one final problem for Steinberger's Directive and other coherence norms. The problem is that factoring the reasons imputed by a coherence norm into your own deliberations seems to amount to a kind of double counting. Alex Worsnip puts the problem quite nicely:

If considerations of coherence constituted reasons, then when I believed that my evidence decisively supports believing p, it would seem to me that I have two reasons to believe p. First, the fact that my evidence decisively supports believing p (or whatever fact or facts make it the case that it decisively supports believing p). And second, the fact that, since I believe that my evidence decisively supports believing p, believing p would make me coherent (or, bring me to satisfy a requirement of structural rationality). But that is not, in fact, how things seem in situations like this one. In situations like this one, I won't treat this coherence-constituted consideration as a second reason to believe p, in addition to that provided by my evidence. From the first-personal point of view, coherence-based reasons seem to be "superfluous"; they don't "add anything... if I did, in my deliberation, treat the fact that believing p would make me coherent as a second, additional reason to believe p, I would seem to be making a mistake. At least assuming that I'm right that my evidence does decisively support believing p, to say that I have two reasons to believe p – first, the fact that my evidence decisively supports believing p, and second, the fact that, since I believe that

¹³⁹ Kolodny 2007, p. 238.

my evidence decisively supports believing p, believing p would make me coherent – seems to be engaging in an illicit kind of “double-counting”. Again, assuming that the first, evidential consideration is a reason to believe p, this suggests that the second, coherence consideration is not a second, additional reason to believe p.”¹⁴⁰

All together, Kolodny’s arguments and the double counting problem make a strong case for NR-. Perhaps, as Kolodny suggests, Evidentialism is true, and adequately explains our intuitions which push us toward accepting coherence norms like Steinberger’s Directive. If nothing can be said for taking Steinberger’s Directive as a plausible directive beyond Evidentialism, then this would make for a strong argument that logic is not, in fact, normative for belief.

The Case For NR+

As I’ve discussed in the last few sections, there is a strong case that the norms of logic are not reducible to evidential norms, and that perhaps evidential norms are the only true directives. This was the thesis I called NR-. If so, this would repudiate the normativity of logic for *belief*, though not necessarily for reasoning.

While I won’t deny that this case is strong, I think there is a case to be made for NR+: that Steinberger’s Directive is a true directive that coexists with our evidential norms, despite not reducing to them. I’ll now present that case.

What Steinberger’s Directive encapsulates is a defeasible requirement to respect the logical commitments of one’s beliefs. There is something amiss about an agent who knows that

¹⁴⁰ Worsnip 2021 pp. 5-6.

$P_1 \dots P_n \models Q$, but who has mismatched attitudes toward the premises and the conclusion. *Pace* Kolodny, it strikes me that an agent *can* in fact utilize this requirement in deliberation.

Let me explain. Consider the familiar case of a philosopher who encounters a new paradox: the philosopher is shown that premises A, B, and C, which are all highly plausible and well supported by the evidence, entail a conclusion D, which is highly unpalatable and runs counter to the philosopher's intuitions or theoretical commitments. What should (in the first-personal, directive sense) the philosopher do?

Here's an intuitive answer. Granting that the conclusion D is sufficiently unpalatable, the philosopher should immediately suspend judgement on the premises A B and C. If the philosopher then interrogates the reasons she has for holding each of the premises and rejecting the conclusion, she may be able to find a solution to the paradox by identifying which of the premises is least plausible, or by finding a flaw in the argument from A, B, and C to D, or by interrogating her reasons for rejecting D, and perhaps coming up with new evidence that suggests that D isn't so unpalatable after all.

The one thing she can't do, however, is simply retain her beliefs in all three premises while rejecting the conclusion. What's more, the philosopher is perfectly well aware of this: she might even motivate her suspension of judgement of the three plausible premises by thinking to herself, "look, no matter how plausible A, B, and C are, one of them must be false, since they entail D and D just can't be true!" From a deliberative standpoint, ignoring the paradox and retaining the premise beliefs while rejecting the conclusion simply isn't a rationally acceptable option.

But note that the true evidential norms by themselves plausibly can't explain this behavior.¹⁴¹ It's certainly true that, once she realizes that $A, B, C \models D$, her evidence for the conjunction $A \& B \& C$ becomes evidence for D . Plausibly, the reverse is also true: her evidence against D counts against the conjunction of the premises (especially since, we can suppose, she recognizes that if $A, B, C \models D$, $\sim D \models \sim(A \& B \& C)$). But even so, having strong reasons against a conjunction of one's beliefs does not straightforwardly entail having strong reasons to reject any of the conjuncts individually, as the Preface Paradox shows. So the norms of evidence won't necessarily require the philosopher to even *suspend* on any of the premises.

Steinberger's Directive provides an explanation for why the philosopher should suspend judgement on the premises. For even though the philosopher doesn't (at least off the bat), have sufficient evidence to reject any of the premises individually, she is in a position to recognize that at least one of the premises must be false. According to Steinberger's Directive, she will thus have defeasible reasons to not believe all of the premises together. Since there are only three premises in this case, the case is not preface-like, so these reasons will be very strong. Thus, the philosopher's best recourse will be to suspend on the three premises, at least until she has the chance to identify which premise is the (or a) culprit. And the philosopher can motivate doing

¹⁴¹ One might object to the case I'm describing by appealing to certain facts about probability. Suppose evidence is probabilistic. Then, if each of A, B , and C have a very high probability, and jointly entail D , there will be limits to how low D 's probability can be, on account of the fact that there are only three premises. If, for example, the probability of D was 0, then if A, B , and C were all equally likely, they could have at most a probability of .666. There are a few things to say in response. First, it won't take too awfully many premises to make a case like what I describe possible even if evidence is probabilistic: it only takes 10 premises to construct a lottery paradox like case where each premise has a probability of .9 while they jointly have a probability of 0, for example. Second, as I've already argued, evidence is hyperintensional, and therefore evidential support need not obey the probability calculus. Third there are no shortage of actual examples of what I'm describing in philosophy: I suspect every philosopher has had the experience of encountering a paradox which follows from a small number of very plausible premises. If probabilistic construals of evidential support can't accommodate this, so much worse for them.

this by appealing to having reason to respect her own logical commitments (i.e., by appealing to Steinberger's Directive).

So what of the double-counting problem? On reflection, it isn't clear how serious of a problem it is. It's certainly true that, if the philosopher identifies A as the culprit and rejects it, she can't cite restoring coherence as a reason for rejecting A. Rather, her reasons for rejecting A will consist in the entailment fact, the reasons for not rejecting B and C, her reasons for rejecting D, and whatever new evidence she gathers to determine that A is less plausible than the other premises. But this isn't so surprising: perhaps all it means is that the reasons to restore coherence lose their force in the presence of decisive reasons for one of the coherence restoring options. Since reasons are highly context sensitive in nature, this isn't an unreasonable suggestion.

So here is my suggestion for the deliberative role of Steinberger's Directive. Agents have reasons to see to it that their beliefs respect the logical commitments that the agents are aware of. While agents can't appeal to these reasons to bolster their support for propositions which their evidence already supports, they can appeal to them when the evidence reveals that the beliefs jointly entail an unpalatable consequence, but it isn't clear which belief to reject. Of course, if the preface paradox is any guide, these reasons are defeasible, and may not even demand suspension of judgement if the set of premises which leads to the unpalatable conclusion is sufficiently large.¹⁴² Nevertheless, they have a genuine role to play in deliberation in some cases.

If I'm correct, there is thus good reason to plump for NR+ over NR-, despite Kolodny's strong arguments. This means that, strictly speaking, Evidentialism is false. But there is

¹⁴² This of course assumes that we don't simply deny the preface paradox entirely. As I've said, I find preface denial rather attractive, in part because it isn't clear to me that merely adding premises should reduce the force of my reasons to respect my logical commitments. But preface denialism is highly controversial, and I do not wish to make it seem like the case for the normativity of logic crucially hinges on our willingness to reject the preface paradox.

nevertheless good reason for the Evidentialism-favoring intuitions on my view: while evidential norms do not exhaust the realm of the epistemic, if your evidence does not demand that you disrespect your logical commitments, your evidence will solely determine what you ought to believe. It is only when the evidence pushes for incoherence that considerations of logicity ever show their face. And even then, your evidence might win, if preface situations are any guide.

Evidentialism and The Normativity of Logic for Reasoning

Before wrapping up, it will be worth briefly discussing whether there is any reason to fear a conflict between the norms of logic on *reasoning* that I defended in Part 4 and Evidentialism. As I've already discussed, on my account, the rules of deductive logic will set the permissible rules for deductive reasoning. The rules of reasoning are thus rules for forming beliefs in *support statements*. These support statements, in turn, can be used as reasons to believe (or disbelieve) other propositions.

Fortunately, while my account is certainly compatible with Steinberger's Directive, it doesn't in anyway hinge on it. Nor is it incompatible with Evidentialism. So even if my arguments for NR+ over NR- fail, there is still good reason to think that logic is normative for reasoning.

On my account, a deductively inferred support statement (a belief whose content has the form $P_1 \dots P_n \models Q$) can count, given suitable circumstances, as a reason for belief in Q. But this need not conflict at all with Evidentialism, because there is no reason to think that said support statement can't be *evidence* for Q. In fact, if as I've been stressing, evidence is hyperintensional, then beliefs in support statements are an extremely important kind of evidence, as they facilitate

the transmission of one's evidence for $P_1 \dots P_n$ to Q . Support statements *enable* premises to lend support to conclusions which follow from them. To be sure, support statements are somewhat peculiar in that they don't by themselves support their conclusions: learning that $P_1 \dots P_n \Rightarrow Q$ doesn't by itself give any reasons to believe or disbelieve Q . One further needs evidence for the premises in order to have evidence for the conclusion. But this isn't so strange: on any plausible account of evidence, there will be propositions which play this sort of enabling role.

To see this, suppose you are one of many candidates for a job, and have the last interview. While in the waiting room, you see a rival candidate wearing a bright red tie, and after commenting on it, you learn that the candidate's name is Steve. After your interview, you overhear the interviewer telling his assistant to inform the man in the red tie that he got the job. On any plausible view, the evidence you receive from overhearing this comment *enables* your prior perceptual evidence of Steve wearing the red tie to become evidence that Steve got the job. That is, your evidence that Steve got the job consists in your memory of Steve wearing the red tie and your perceptual evidence of the interviewer's comment. Either of these alone gives you no evidence whatsoever that Steve got the job; only in conjunction do they support this conclusion.

What this shows is that there is nothing overly unusual about evidence playing an enabling role, in the way that support statements do, on my view. So we can happily accept my account of reasoning, and the normativity of logic for reasoning as I defend it, even if Evidentialism is true.

Final Conclusion

In this thesis, I've investigated the normativity of logic for belief and for reasoning. After setting the stage regarding what I take logic to be in Part 1, and giving a primer on epistemic normativity in Part 2, I turned to the normativity of logic for belief. In part 3, I drew from some recent work by Florian Steinberger, and argued that logic is normative for belief, in that there are true bridge principles between the facts of logic and various kinds of norms on belief. In part 4, I turned to the normativity of logic for reasoning. I presented an account of reasoning inspired by the views of Paul Boghossian and Pavel Tichý, in which reasoning is a rule-following affair. I argued that logic is normative for reasoning in a very robust sense: logic sets constitutive norms for deductive reasoning, which are satisfied when an agent reasons according to the rules of logic which they are justified in applying.

In Part 5, I investigated how these norms fit into an overall theory of the epistemic. I argued that Evidentialism is best construed as setting epistemic directives for belief, and then investigated how Steinberger's Directive relates to Evidentialism. I argued that it is hard to see how Steinberger's Directive could be reducible to Evidentialism, and that Evidentialism certainly isn't reducible to Steinberger's Directive. In light of the overwhelming plausibility that there are at least some true evidential directives, this either means that Evidentialism is true and Steinberger's Directive is false (a thesis I dubbed NR-), or means that Steinberger's Directive is true, and coexists with the true evidential norms (a thesis I dubbed NR+). I presented a very strong extended argument for NR- stemming from Niko Kolodny, and then argued that, despite this argument's strength, there was good reason to opt for NR+. Finally, I discussed how the norms of logic for *reasoning* interact with Evidentialism, and argued that they were fully

compatible with it, and even helped explain how evidence is transmitted in light of the hyperintensionality of evidence.

The final upshots are the following. First, the normativity of logic *for belief* appears to be incompatible with Evidentialism (though is compatible with Weak Evidentialism). Second, if I'm correct, we should nevertheless accept that logic is normative for belief, with the primary logical norm being Steinberger's Directive. Third, whether we accept the normativity of logic for belief or reject it out of allegiance to Evidentialism, if I'm correct, logic is normative *for reasoning*, in a very robust sense.

Bibliography

- Balcerak Jackson, Magdalena and Balcerak Jackson, Brendan. 2013. "Reasoning as a Source of Justification". *Philosophical Studies* 164: 113-126.
- Berker, Selim. 2013a. "Epistemic Teleology and the Separateness of Propositions". *The Philosophical Review* 122 (3): 337-393.
- Berker, Selim. 2013b. "The Rejection of Epistemic Consequentialism". *Philosophical Issues* 23: 363-387.
- Berker, Selim. 2018. "A Combinatorial Argument Against Practical Reasons for Belief". *Analytic Philosophy* 59 (4): 427-470.
- Boghossian, Paul. 2003. "Blind Reasoning". *Proceedings of the Aristotelean Society* 77: 225-248.
- Boghossian, Paul. 2014. "What is Inference?". *Philosophical Studies* 169: 1-18.
- Broome, John. 2000. "Normative Requirements". In J. Dancy, (Ed.), *Normativity*. Oxford: Oxford University Press.
- Broome, John. 2013. *Rationality Through Reasoning*. Oxford: Wiley Blackwell.
- Broome, John. 2014. "Comments on Boghossian". *Philosophical Studies* 169: 19-25.
- Christensen, David. 2004. *Putting Logic in its Place: Formal Constraints on Rational Belief*. Oxford: Clarendon Press.
- Cohen, Stewart. 2016. "Theorizing About the Epistemic". *Inquiry* 59 (7-8): 838-857.

- Cook, Roy. 2010. "Let a Thousand Flowers Bloom: A Tour of Logical Pluralism". *Philosophy Compass* 5 (6): 492-504.
- Drayson, Zoe. 2014. "The Personal/Subpersonal Distinction". *Philosophy Compass* 9 (5): 338-346.
- Feldman, Richard. 2000. "The Ethics of Belief". *Philosophy and Phenomenological Research* 60 (3): 667-695.
- Feldman, Richard. 2003. *Epistemology*. Upper Saddle River, N.J.: Prentice Hall.
- Field, Hartry. 2009. "What is the Normative Role of Logic?". *Proceedings of the Aristotelian Society* 83: 251-268.
- Fine, Kit. 1985. *Reasoning with Arbitrary Objects*. Oxford: Basil Blackwell.
- Forbes, Graeme. 1993. "But a Was Arbitrary...". *Philosophical Topics* 21 (2): 21-34.
- Forbes, Graeme. 1994. *Modern Logic: A Text in Elementary Symbolic Logic*. New York: Oxford University Press.
- Frege, Gottlob. 1906. "Über die Grundlagen der Geometrie II". *Jahresberichte der Deutschen Mathematiker-Vereinigung* 15: 58-77.
- Frege, Gottlob. 1979. *Posthumous Writings*. Hans Hermes, Friedrich Kambartel and Friedrich Kaulbach (eds.). Oxford: Blackwell.
- Friedman, Jane. 2017. "Why Suspend Judging?". *Nous* 51 (2): 302-326.
- Gentzen, Gerhard. 1934. "Untersuchungen über das logische Schliessen". *Mathematische Zeitschrift* 39: 176-210, 405-431.

Gentzen, Gerhard. 1936. „Die Widerspruchsfreiheit der reinen Zahlentheorie“. *Mathematische Annalen* 112: 493-565.

Hacking, Ian. 1979. “What is Logic?” *The Journal of Philosophy* 76 (6): 285-319.

Harman, Gilbert. 1986. *Change in View: Principles of Reasoning*. Cambridge: MIT Press.

Hedman, Shawn. 2004. *A First Course in Logic*. New York: Oxford University Press.

Hlobil, Ulf. 2014. “Against Boghossian, Broome, and Wright on Inference”. *Philosophical Studies* 167: 419–429.

Horty, John. 2012. *Reasons as Defaults*. New York: Oxford University Press.

Humberstone, Lloyd. 2000. “The Revival of Rejective Negation”. *Journal of Philosophical Logic* 29: 331–381.

Hurley, Patrick J. *A Concise Introduction to Logic*. Stamford: Cengage Learning.

Kolodny, Niko. 2007. “How Does Coherence Matter?”. *Proceedings of the Aristotelian Society* 107 (3): 229-263.

MacFarlane, John. 2004. “In What Sense (If Any) Is Logic Normative for Thought?”.
Unpublished: 1-25.

MacFarlane, John. "Logical Constants". In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2017 Edition).

<https://plato.stanford.edu/archives/win2017/entries/logical-constants/>.

McHugh, Conor and Way, Jonathan. 2016. “Against the Taking Condition”. *Philosophical Issues* 26: 314 – 331.

- McHugh, Conor and Way, Jonathan. 2017. "What is Reasoning?". *Mind* 127 (505): 167-196.
- McSweeney, Michaela Markham. 2019. "Logical Realism and the Metaphysics of Logic".
Philosophy Compass 14 (1): 1-10.
- Milne, Peter. 2009. "What is the Normative Role of Logic?". *Proceedings of the Aristotelian Society* 83: 269–298.
- Murzi, Julien and Steinberger, Florian. 2017. "Inferentialism". In B. Hale, C. Wright, and A. Miller (Ed.), *A Companion to the Philosophy of Language*. 2nd Edition. 197-224.
- Neta, Ram. 2019. "The Basing Relation". *Philosophical Review* 128 (2): 179-217.
- Pezlar, Ivo. 2013. "Tichy's Two-Dimensional Conception of Inference". *Organon F* 20 (2): 54-65.
- Peacocke, Christopher. 1987. "Understanding Logical Constants: A Realist's Account".
Proceedings of the British Academy 73: 153-200.
- Prior, Arthur. 1960. "The Runabout Inference-Ticket". *Analysis* 21 (2): 38-39.
- Restall, Greg. Forthcoming. "Speech Acts and the Quest for a Natural Account of Classical Proof". Unpublished.
- Rinard, Susanna. 2019. "Equal Treatment for Belief". *Philosophical Studies* 176: 1923–1950.
- Richard, Mark. 2019. "Is Reasoning a Form of Agency?". In M. Balcerak Jackson and B. Balcerak Jackson (Eds.), *Reasoning: New Essays on Theoretical and Practical Thinking*. Oxford: Oxford University Press.
- Rumfitt, Ian. 2000. "'Yes' and 'No'". *Mind* 109 (436): 781-823.

- Russell, Gillian. "Logical Pluralism". In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2021 Edition).
<https://plato.stanford.edu/archives/sum2021/entries/logical-pluralism/>.
- Sainsbury, Richard M. 2002. "What Logic Should We Think With?". In A. O'Hear (Ed.), *Logic, Thought and Language*. Cambridge: Cambridge University Press.
- Searle, John. 1969. *Speech Acts*. Cambridge: Cambridge University Press.
- Sher, Gila. 2011. "Is Logic in the Mind or in the World?". *Synthese* 181: 353-365.
- Sider, Ted. 2010. *Logic for Philosophy*. New York: Oxford University Press.
- Silva, Paul and Oliveira, Luis. 2020. "Propositional Justification and Doxastic Justification". *The Routledge Handbook of Evidence*.
- Stapleford, Scott. 2016. "Epistemic Value Monism and the Swamping Problem". *Ratio* 29: 283-297.
- Steinberger, Florian. 2011. "What Harmony Could and Could Not Be". *Australasian Journal of Philosophy* 89 (4): 617-639.
- Steinberger, Florian. 2016. "Explosion and the Normativity of Logic". *Mind* 125: 385-419.
- Steinberger, Florian. 2017a. "Frege and Carnap on the Normativity of Logic". *Synthese* 194 (1): 143-162.
- Steinberger, Florian. 2019a. "Three Ways in Which Logic Might Be Normative". *The Journal of Philosophy* 96 (1): 5-31.

- Steinberger, Florian. 2019b. “Consequence and Normative Guidance”. *Philosophy and Phenomenological Research* 98 (2): 306-328.
- Steinberger, Florian. 2019c. “Logical Pluralism and Logical Normativity”. *Philosopher’s Imprint* 19 (12): 1-19.
- Streumer, Bart. 2007. “Reasons and Entailment”. *Erkenntnis* 66: 353–374.
- Tarski, Alfred. 1966. “What are Logical Notions?” *History and Philosophy of Logic* 7: 143–154.
- Tichý, Pavel. 1988. *The Foundations of Frege’s Logic*. Berlin, New York: De Gruyter.
- Turri, John. 2010. “On the Relationship Between Propositional and Doxastic Justification”. *Philosophy and Phenomenological Research* 80 (2): 312-326
- Wedgwood, Ralph. 2006. “The Normative Force of Reasoning”. *Nous* 40 (4): 660–686.
- Worsnip, Alex. 2018. “The Conflict of Evidence and Coherence”. *Philosophy and Phenomenological Research* 96 (1): 3-44.
- Worsnip, Alex. 2021. “Making Space for the Normativity of Coherence”. Forthcoming in *Nous*: 1-23.
- Valaris, Markos. 2017. “What Reasoning Might Be”. *Synthese* 194: 2007–2024.