Inference and the Normativity of Logic

James R. Shaw [†]

"Logic" is a term of art with imprecise application. But there is one reasonable, and I would argue central, construal of logic which has been absent from the most influential recent discussions of logical normativity. This construal takes logic to be the linguistically mediated study of good deductive inference where, critically, deductive inference is construed as special kind of contentful mental activity. On this view, logic tracks a fundamentally *evaluative* status (goodness) which governs a mental *act or process* (inference). The aim of this paper is to argue that these two features, and especially their interaction, provide the keys to a proper understanding of logic's normative role.

I begin in §1 by briefly sketching the view of logic I have in mind. In §2, I provide a critical overview of the literature on bridge principles, which typically frame logical normativity using deontic (rather than evaluative) language that places constraints on combinations of attitude states like beliefs or credences (rather than on a mental act). In §3, I note that a view that instead treats logic as tracking an evaluative status governing mental acts of inference simply sidesteps the ever growing array of obstacles to the views of §2, and additionally reveals that the seemingly diverse array of such obstacles can in fact be traced to a simple, common source. Finally, in §4, I step back to evaluate how the view of the normativity of logic I put on offer engages with the theoretical motivations of the authors I argue against, focusing on the question of what role logic plays within our accounts of good reasoning.

1   Inference, Inferential Goodness, and Logic

Let me begin by giving an account of inference, and inferential goodness, before explaining how we can understand logic as investigating it. My discussion of inference will be very brief, since my primary aim is to highlight how even the general shape of such an account can impact our conception of logical normativity.

What is an inference? A necessary condition will help to get us started: inference is a process in which one acceptance state is generated on the basis of others. I believe that Joppa is north or south, and believe it's not north; I infer that Joppa is south. In so doing, I come to a new belief on the basis of two others. But I could equally *suppose* that Joppa was north or south, suppose further it was not north, and infer 'under supposition' that it was south. In so doing, I merely suppose something new on the basis of two suppositions. Likewise I seem able to infer my conclusion from imagined premises, if imagination is different from supposition.

Belief, supposition, and imagination are what STALNAKER (1984) calls states of *acceptance*, offering as a diagnostic that these are states we call "correct" just in case their contents are true. One believes, supposes, or imagines correctly just in case one believes, supposes, or imagines what is the case.[1] Acceptance states contrast with preferential states like desire or hope, which are not said to be correct if what is preferred transpires. And they contrast with inquisitive states (like a state of wondering whether Joppa is north) which do not even take a truth-evaluable object. Notably, these latter states cannot participate in inference: one cannot infer from, or to, a state of desire or hope or wonder.

An alternative, equally important understanding of acceptance states is as *information bearing* states—those in the business of representing how the world is or might be. The intentional structure of total mental states of this kind are fruitfully modeled by something like sets of worlds or sets of propositions. Not so for preferential states which represent how a world is preferred to be. These are familiarly construed not as a collection of worlds or propositions, but as a *ranking* of worlds or propositions.[2] Not so for an inquisitive state, like a state of wondering, which is better modeled by something like a *partition* on worlds, or sets of sets of alternative propositions.[3]

---

[1]   I will not be presuming, nor do I suspect, that the use of "correct" here tracks anything interestingly normative.

[2]   See BOLINGER (1968), STALNAKER (1984), FARKAS (1985), and HEIM (1992).

[3]   See FRIEDMAN (2013) on inquisitive attitudes and GROENENDIJK & STOKHOF (1997) for a discussion

What this tells us is that inference is an act bridging all and only mental states with an inherently information-bearing, correctness-governed character. This helpfully constrains our understanding of inference if, as I propose, we should aim to understand what inference is along broadly functionalist lines by what it does. We should construe inference as a mental activity which has some useful role to play in application to information bearing states, and only makes sense in application to them.

In what follows it will be helpful to fix ideas with a working proposal, even though any such view is bound to be controversial. As I proceed, I will try to flag if a controversial element of the proposal is doing integral work. My (largely unoriginal) working proposal is as follows:

> *An inference is a mental act whose proper function is to appreciably generate new mental states on the basis of old ones in a reliably correctness-preserving way.*

Let me unpack some of the terminology.

> *…a mental act…*: Following standard views, I take inference not to be a particular kind of mental state (like belief), but an act or activity that mediates between such states.[4]

> *…proper function…*: In speaking of inference's function, I appeal to a broadly functionalist tradition in the philosophy of mind according to which mental phenomena are individuated by their functional role in a cognitive system. More specifically, in speaking of *proper* functions, I appeal to a version of functionalism incorporating teleological elements, which is how a form of normativity will enter the picture. On this functionalist tradition, we conceptualize certain entities, including biological entities like hearts or kidneys as being *for* certain ends or purposes, like pumping or filtering blood. In being conceptualized with purposes, such entities establish a standard by which they can be said to be good instances of their kind.[5] For example, a heart is good (*qua* heart) provided it pumps

---

of the supporting literature on the semantics of interrogatives.

[4] See BUCKAREFF (2005), GIBBONS (2009), HIERONYMI (2009), MELE (2009), PEACOCKE (2008). For dissent see STRAWSON (2003), SETIYA (2013). If one doubts inference is an act, many of the claims about logic to follow can plausibly be defended on the much weaker assumption that inference is a process or event. What is important is that it is not simply a state.

[5] See ZIFF (1960), FINLAY (2004, 2014), THOMSON (2008). The broad set of ideas here of course goes back at least as far as Aristotle. The application of them to the psychological sphere has its earliest

blood well, bad (*qua* heart) if it pumps poorly. I propose that when we call an inference good we draw on a functional conceptualization of inference, associating it with a purpose establishing its attendant standard of goodness. That purpose is spelled out by the three further features explained below.

…correctness preserving…: Correctness for an acceptance state is, we have noted, truth. The conditions that matter to the correctness of such a state are, therefore, those that matter to truth. These, familiarly, are encapsulated in the notion of a *possible world*. So correctness is world-relative. Say that an acceptance state transition is correctness-preserving at a world *w* just in case the based attitudes are correct at *w* if the basing attitudes are. The current proposal is that inference's proper function is to effect a correctness-preserving operation of some kind. This supposition has a key explanatory virtue: it explains why all and only correctness-bearing states can participate in inference. Only they have the correctness-governed character which inference works to preserve.

…appreciably…: A conscious, successful inference is one whose distinctive correctness-preservation is somehow recognized by the inferrer.[6] When one consciously infers, one somehow 'takes' one's inference to be well-performed. Though it is notoriously difficult to specify what this appreciation comes to,[7] there are good reasons to consider the appreciation condition to be a rational commitment of the inferrer.[8] On the whole, I will remain neutral on what appreciation involves.

…reliably…: We've seen that correctness (or truth) preservation is world-relative. So which worlds matter to inference well performed?[9] A tempting

---

6   detailed development in Millikan (1984), though I don't mean to plump for any particular (for example evolutionary) understanding of these teleological conceptions here—a general framework neutral on these issues will suffice.

6   Cf. Locke (1690/1979), Frege (1879?/1979), Russell (1920/1988), Stroud (1979), Thomson (1965), Sainsbury (2002), MacFarlane (ms/2004), and Field (2009). I allow for the existence of unconscious inferences. These needn't be appreciated, but seemingly need to be appreci*able* (to the inferrer) to do their job—hence the condition as stated.

7   See in this regard the recent debate among Boghossian (2014), Wright (2014), Broome (2014), and Hlobil (2014).

8   As argued by Hlobil (2019), consciously drawing an inference while believing it not to be good leads to a kind of Moorean rational incoherence.

9   Also which *kind* of worlds matter? I take good inference to require correctness preservation at meta-

option should be summarily dismissed: inference's function sometimes requires more, and sometimes less, than *actual* correctness-preservation. It may require more, because all inferences that start from suppositions of actual falsehoods preserve actual correctness of contents trivially. But obviously not all such inferences are well-performed. And it may require less, because ampliative inference may be well-performed even if it does not preserve actual correctness (because the actual world is an unusual or unlikely case). I think counterfactual inference helps reveal most clearly what inference, *qua* inference, aims to effect: it aims to preserve correctness across a 'safe' range of cases *compatible with the information contained in the starting acceptance states* on which inference operates. Ampliative inference may reliably preserve correctness in this way, even if it leads from a correct state to an incorrect state (since the actual case may not have been among the 'safe' cases). And for suppositional inference to be reliably correctness-preserving may require more than being actually correctness-preserving, since many worlds besides the actual one may be compatible with the starting suppositional state. Indeed, the actual world may not even be compatible with it.

This proposal is obviously skeletal, and should remain so, since some details of the metaphysics of inference aren't relevant for my investigation here. But one aspect of the proposal will need additional clarification: What does reliability come to? What does it take for a range of worlds compatible with the information in an information bearing state to be 'safe' for an inferential transition?

This is a vexing question, because of how tricky it is to account for the conditions on good ampliative inference. Fortunately I can sidestep the hard part of the question: my ultimate interest is in deductive logic, and so with deductive inference—so from here on out, talk of inference will always be talk of deductive inference.[10] And the safety conditions on deductive inference are simple: they clearly involve 'maximal safety.' In other words, deductive inference aims at the preservation of correctness (or truth) at *all* worlds compatible with starting acceptance states.

---

physical possibilities, but will not defend this view here. I do not see that the arguments of this paper turn on this issue.

[10] Is deductive inference a special *kind of inference* which fails in its function even if it is good by ampliative standards, or is deductive-correctness merely a *kind of correctness*, not corresponding to any distinct kind of inference, but rather carving out one of several ways in which inference (*simpliciter*) can count as correctly performed? I don't see much in what follows as hanging on this question, so I won't settle on an answer. For simplicity and consistency, I will speak of deductive inference in the former sense.

Note that the worlds incompatible with an acceptance state are precisely those at which it is incorrect. This means that the condition that good inference preserve correctness at all worlds compatible with starting acceptance states is equivalent to a simpler condition: preserving correctness at all worlds. (I've formulated this proper function of deductive inference in terms of correctness, but I could have equally formulated it in terms of truth-conditional information: deductive inference aims at a total information-preserving transition between acceptance states. An inference should move one to an acceptance state that 'rules out' no more worlds than the states with which one began. This is, again, readily seen to be equivalent to preservation of truth at all worlds.)

This concludes my sketch of inference and its attendant standard of goodness. To summarize: once we have information-bearing, correctness-governed states that represent how the world is or merely might be, it is helpful to have a kind of mental activity whose proper function is to appreciably and safely extract implicit informational commitments of those states. Inference is the mental activity geared to effect such an operation. Inference counts as good (*qua* inference) when it successfully fulfills this role—its proper function.

Now, how can deductive logic contribute to an investigation of good deductive inference, so construed? Recall that there are two features of good deductive inference: (a) they preserve truth at all possible worlds and (b) this fact is appreciable to the inferrer.

As many have noted, the second condition here is not easily subject to systematic investigation, owing to its psychological variability. A Ramanujan may inferentially flit through the space of mathematical possibility, correctly seeing 'obvious' steps in ways that simply baffle the ordinary reasoner. On the other end of the spectrum, as CARROLL (1895) effectively notes, it seems possible for the 'phenomenally obtuse' to be unable to see even the most elementary of acceptable inferential transitions. What this means is that appreciability is not obviously subject to systematization without arbitrary stipulations.

By contrast, the first condition on good deductive inference, (a), ends up being a simple modal property of mental content, completely divorced from contingent psychology. We could not hope for a condition more amenable to systematization. We can investigate modal properties of mental content in the customary way: presume that such contents are expressed by the declarative sentences of a (formal or natural) language. This allows us to exploit the compositional structure of such sentences in sys-

tematizing necessary truth-preserving relations among such contents—thereby, I claim, investigating a necessary condition on good deductive inference.

To engage in this task is so far to engage in the task of investigating general *entailment relations* in compositional semantic theorizing. Such entailments would subsume necessary truth-preserving relations that are not traditionally conceived of as logical. So-called lexical entailments (like the entailment from something's being a vixen to its being a fox) may be of this form.[11]

Accordingly, to get from here to logical inquiry proper, we make a broadly familiar move: we note that sometimes sentence transitions express contents that can be seen to necessarily preserve truth 'in virtue' of some restricted set of their linguistic properties.[12] For example, an entailment between sentences may be guaranteed merely under the assumption that its predicate denotations belong to some general class of predicate denotations, rather than bearing the specific denotations they do. If so, a logic may abstract from the details of particular predicate denotations in tracking entailment relations, instead holding constant the denotations of certain special 'logical' vocabulary.

A *logic*, then, investigates relations of necessary truth-preservation among sentences' assertoric contents in virtue of a particular, hopefully revealing and formalizable subset of their linguistic properties.

So much for logic and its role in studying inference. Obviously much more would be needed to clarify, refine, and defend this conception of logic. But this is not my aim here. It turns out that even with this bare sketch, we have all the elements we will need to investigate the sense in which logic, understood as I have, is normative.

2   Bridge Principles and their Discontents

The normative scope of logic, on the foregoing view is as follows: Logic…

(i)   …*non-exhaustively* tracks a …

(ii)   …*necessary* but *insufficient* condition on …

(iii)   …an *evaluative* normative status governing…

(iv)   …*acts* (and not states).

---

[11]   See Glanzberg (2015) for a helpful discussion of such entailments and their relation to logic.

[12]   Cf. Sánchez-Miguel (1992), Etchemendy (2008).

Logic tracks a necessary condition (necessary truth-preservation) on the evaluative status (goodness) of an act (inference). It does so non-exhaustively (ignoring the goodness connected with, e.g., lexical entailments). And the necessary condition it tracks is insufficient for the evaluative status (since logic sets aside psychologically variable appreciability requirements on inferential goodness).

How might logic and its normative scope, so understood, engage with existing debates? To begin, the view would have immediate and transformative implications for a burgeoning program in philosophical logic: the task of finding normative 'bridge principles' for logic. The terminology derives from MacFarlane (ms/2004), but the project owes its life to a kind of skeptical challenge dating back to Harman (1984, 1986).

In this section, I give a brief and opinionated survey of the state-of-play in providing bridge principles. The discussion will be quick, because my goal is less to do justice to the details of this program than it is to explain how the view of §1 can fruitfully intervene in the debates.

In the course of asking how logic relates to good reasoning broadly construed (a question I will take up in §4), Harman suggests that logic's role in reasoning would be cashed out with something like the following forms.

> LOGICAL IMPLICATION PRINCIPLE: The fact that one's view implies P is a reason to accept P.

> LOGICAL INCONSISTENCY PRINCIPLE: Logical inconsistency is to be avoided.

MacFarlane (for different theoretical purposes) generalizes the idea by introducing a schematic form for such principles, dubbing their instances *bridge principles*. These forms link facts about logical entailment and normative claims about attitude states roughly as follows:

> If $A, B \models C$ then [normative claim about believing A, B, and C]

MacFarlane notes that we can view the space of such bridge principles by varying, as he puts it, the "type of deontic operator" (he considers: obligations, permissions, reasons), its "polarity" (whether it prescribes believing, or not disbelieving), the scope of the deontic operator, and whether the antecedent of the bridge principle must be known for the principle to hold. After surveying 36 formulations, MacFarlane tentatively seems to endorse the following.

(wo-) If $A, B \models C$ then you ought to see to it that if you believe A and you believe B, you do not disbelieve C.

(wr+) If $A, B \models C$ then you have reason to see to it that if you believe A and you believe B, you believe C.

The proposal has spawned a number of competitors (or supplements, depending on the theorist). FIELD (2009) extends bridge principles to apply to credences, eventually proposing a generalization of the following principle:

(D*) If it's obvious that $A_1, \ldots, A_n \models B$, then one ought to impose the constraint that $P(B)$ is to be at least $P(A_1) + \ldots + P(A_n) - (n-1)$, in any circumstance where $A_1, \ldots, A_n$ and B are in question.

STEINBERGER (2019a), who focuses on finding bridge principles that could encapsulate logic's guidance of good reasoning, makes the following proposal:

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

When Harman formulated his original two principles, he did so precisely to highlight their susceptibility to counterexample. Indeed, he raised a barrage of such worries, since expanded by other philosophers. In brief, these include:[13]

BACKTRACKING: A recognized entailment can be a reason to abandon one's starting beliefs if one has sufficient evidence against its consequences.

BOOTSTRAPPING: Although in most logics $p \models p$, it does not appear that simply believing p guarantees one has reason to believe p.

CLUTTER: It is an irrational waste of our finite cognitive resources to needlessly clutter our minds with irrelevant consequences of our current beliefs.

PARADOXES: It appears rational to respond to some unresolved paradoxical or puzzling situations by maintaining logically inconsistent beliefs and managing

---

[13] See HARMAN (1984, 1986), BROOME (1999) MACFARLANE (ms/2004), STEINBERGER (2019a) for more extended discussion.

them responsibly in the interim. Cases include (though may not be limited to): the preface paradox, in which one can reasonably believe of any large class of one's beliefs that one of them is false on broadly probabilistic grounds; the liar or Sorites paradox, in which a small set of highly cherished principles seem to lead directly to contradiction; and cases of conscious anti-expertise, where one recognizes that some proposition p is true just in case one fails to believe p (or know it, or have high credence in it).

Excessive Demands: We are not irrational for failing to believe logical entailments of our beliefs that are sufficiently hard to recognize.

All these problems apply pressure to Harman's Logical Implication Principle, and Paradoxes in its many forms seems to undermine even the Logical Inconsistency Principle.

Harman's reaction was to embrace the defeasibility of his principles. But the other philosophers I've mentioned were instead emboldened to revamp them. The revamped principles may improve on Harman's in some ways. But often enough, it is unclear how they even escape Harman's original worries.

Take MacFarlane's (wo-).

(wo-) If $A, B \models C$ then you ought to see to it that if you believe A and you believe B, you do not disbelieve C.

This is essentially a carefully worded version of Harman's Logical Inconsistency Principle, which ran into trouble with Paradoxes. How does the new principle avoid those worries?

MacFarlane focuses on the preface paradox, discussing two strategies for safeguarding (wo-). The first is to admit the existence of conflicting rational norms: in the preface paradox, while it may be true that one ought to maintain one's current inconsistent set of beliefs (since they are duly responsive to the evidence), perhaps it is *also* true that one ought to revise those beliefs to render them inconsistent. Of course, one can't fulfill both these obligations. But, as MacFarlane notes, the existence of conflicting norms is something we are familiar with from other domains (notably the legal).

There are two ways of understanding this proposal. On the first, the conflicting norms are both norms of subjective rationality; on the second, at least one norm (probably the logical norm) is a norm of objective rationality.

The option on which the norms are subjective simply appears false. Not, or not merely, because it posits conflicting subjective rational obligations. But simply because one half of the obligations it appeals to—that one is under a subjective rational obligation to change one's beliefs to render them consistent—do not intuitively obtain. The only motivation for positing the obligation seems to be the *ad hoc* grounds that it would make our logical norm general. Even barring this worry, it is hard to overstate the cost of the first option. All of us are, on brief reflection, in the circumstances of a preface-paradox with respect to some large class of our beliefs. The proposal on the table is that virtually every single reflective agent that has ever existed was *inescapably* subjectively irrational. It is not merely that the proposal jettisons intuitive subjective epistemic ought-implies-can principles (which many theorists, myself included, would already view as a high cost). Rather it makes subjective rationality so demanding as to be transparently unobtainable. What is more, we are entertaining this cost in an effort to secure the indefeasible subjective rational force of logic. It is surely a pyrrhic victory when we have secured the inescapability of logic's norms by rendering norms of their type practically impossible to follow.

The second option, on which logical norms are more objective, fares better in all these respects. On this view, one 'ought' to change one's beliefs in a preface-paradox case only in a sense that somehow takes into account features that goes beyond one's current epistemic limitations. But problems arise when we ask in what particular way the norm would be more objective. It cannot be that one ought to change one's beliefs if only one could *reason better*. A characteristic feature of the preface paradox is that it persists even in the face of completely idealized capacities for reasoning. This seems to leave only one alternative: that the norm is tracking what one ought to believe if only one had *more evidence*. The claim that one 'ought' to change one's beliefs given sufficient added information is plausible. What seems implausible is that this is the form that the norms of logic take. It may be an open question whether logic has normative force. But its having force only for those fortunate few with enough evidence to skirt all preface paradox (again: no actual ordinary agent will qualify) seems about as good as having no force at all.

MacFarlane considers a second approach to the preface paradox that makes no use of conflicting norms. The strategy begins by insisting, counterintuitively, that one ought to render one's beliefs consistent in preface paradox cases. But we soften this counterintuitive character by noting one way that almost anyone could in fact satisfy

such norms: by seeking further evidence. Though one currently has inconsistent beliefs, it is plausible to think that there is evidence that would resolve it. Perhaps logic instructs you to try to find it.

I worry that this suggestion confuses epistemic with practical normativity: I'm not sure practical action like reading a book is ever a way to satisfy epistemic requirements of the form we were seeking with bridge principles. But there is a much simpler concern. The strategy MacFarlane proposes isn't general. There are cases of the preface paradox where no further evidence exists, and one knows this. Suppose you have written a lengthy book on sea turtles, but you are the sole individual rescued by aliens shortly before the Earth's destruction in the crossfire of intergalactic war. Earth, its inhabitants, and the legacy of its sea turtles have been vaporized. Even if you believe you have some mistaken belief about sea turtles, there is no obligation, subjective or objective, to seek out further evidence to rectify the situation. That would be a tremendous waste of time, given that you know there is no such evidence remaining.

So far I've been arguing that MacFarlane's principle runs headlong in the very kind of concerns Harman raised for it, in spite of the inventive ways MacFarlane suggests to avoid them. And this after considering only *one* of the many kinds of tricky epistemic scenarios falling under the heading of PARADOX. This is important to bear in mind when we turn to consider, say, Field's bridge principle applying to credences.

(D*) If it's obvious that $A_1, \ldots, A_n \models B$, then one ought to impose the constraint that $P(B)$ is to be at least $P(A_1) + \ldots + P(A_n) - (n-1)$, in any circumstance where $A_1, \ldots, A_n$ and B are in question.

Field's proposal has the obvious virtue of handling the preface paradox neatly, and in a familiar way: by having a slightly reduced degree of confidence in each member of a large set of propositions, our logical norms can license a very low credence in their conjunction. But it is not clear it can cope will other problems raised by strange and puzzling epistemic circumstances.[14]

---

[14] In particular, though I don't have the space to discuss the matter here, I suspect Field's norm sometimes runs into problems with cases of anti-expertise, especially in light of work by CAIE (2013) which extends the challenges they raise to the credal setting. Caie argues that provided devices of self-reference are available to generate cases of anti-expertise, an agent who is somewhat sensitive to her own credal states and aware of her anti-expertise cannot have credences satisfying the probability axioms. Field's logical norm doesn't strictly speaking require one's credences to satisfy the probability axioms. But neither does Caie's impossibility proof strictly speaking require the probability axioms

But let me set this concern aside for now in favor of a simpler issue. Suppose Field's principle is true and general. What is it illuminating about the normative role of logic in particular? Note, for example, that if we take out the logical relation of entailment, and simply replace it with a conditional, we get a principle which looks just as plausible, and apparently quite a bit more general.[15],[16]

($\overline{D^*}$) If it's obvious that if $A_1 \wedge \ldots \wedge A_n$, then B, then one ought to impose the constraint that $P(B)$ is to be at least $P(A_1) + \ldots + P(A_n) - (n-1)$, in any circumstance where $A_1, \ldots, A_n$ and B are in question.

But this principle doesn't seem to have anything special to do with logical consequence. We can see this even more clearly when we compare what Field's principle is telling us about the normativity of logical truth, (D^-), and comparing it with a principle that replaces talk of logical truth with talk of simple truth, $(\overline{D^-})$.

(D^-) If it's obvious that $\models B$, then one ought to impose the constraint that $P(B)$ is 1.

($\overline{D^-}$) If it's obvious that B is true, then one ought to impose the constraint that $P(B)$ is 1.

Again, the latter principle seems about as plausible, while being significantly more general and more fundamental: if we should assign high credence to obvious logical truths, it seems plausible that this is *because* they are obvious truths.

We should concede that Field's proposal is one possible story about the normativity of logic. On that proposal, logic is normative simply because (actual) truths are, at

---

to get up and running. The dialectical situation is a little complex: depending on what logic one endorses, one may get pressure from Caie's proof to admit the defeasibility of (D^*).

[15] Note that even though we are forced to group the premises into conjunctive form, our principle still gets a non-trivial, and intuitively correct, verdict on the relation between credences in conjunctions and credences in their conjuncts, provided conditionals with the same antecedent and consequent are obviously true.

[16] Field's principle appeals to a constraint that a logical entailment be 'obvious' (in order to avoid troubles from Excessive Demands). It will not be all that important how this epistemic notion is understood beyond that it is supposed to ensure that failure to recognize an obvious logical entailment is not an excuse for contravening logic's demands, while the unobviousness of a violated logical truth can be exculpating. In the principles I use for comparison below, ($\overline{D^*}$) and ($\overline{D^-}$), my ensuing arguments will go through while reading "obvious" in them as strongly as we like, provided the logical principles Field wants to consider 'obvious' in his sense are also 'obvious' in this more general, and stronger sense. As long as this happens, the analog principles I discuss that do not involve logical notions will still be more general and more fundamental.

least when somehow 'obvious'. If something is obviously true, you should probably be confident in it. And if that obvious truth is a truth-functional compound, that has implications for how you confident you should be in its truth-functional constituents. Logic is then normative because, in furnishing us with validities and consequences, it thereby furnishes us with actual truths, including some truths in conditional form.

I don't want to disagree with this story. But I think that if it were *all* to say about the normativity of logic, it would be disappointing. It's customary to treat logical truths as true, and very common to take logical consequences to deliver true conditionals.[17] As long as truths (perhaps the obvious or clear ones) have some role to play in regulating belief, logic will. But this would make logic as distinctively normative as any other domain of inquiry that has some 'obvious' truths in it. And it would mean that the methodology of current theorists is getting things backwards: we should start by figuring out what the normative implications of truth, or recognized truth, are, and derive those for logic as a byproduct.[18]

This concern arises for Field in part because of the rider that logical truths be 'obvious' to avoid certain worries from Excessive Demands. I think a related kind of concern may arise for Steinberger's principle (S).

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

Again, to the extent that (S) is plausible, it is unclear what is special about the logical character of the norm. Norms like $(\bar{S})$ seem about a plausible, and more general.

$(\bar{S})$ If according to S's best estimation at the time, S takes it to be the case that $A_1 \wedge \ldots \wedge A_n \rightarrow C$ and S has reasons to consider or considers C, then S has reasons to (believe C, if S believes all of the A_i).

---

[17] Note that this is not presupposing a deduction theorem, but only that consequence delivers *actual* true conditionals. Still, this too, can be doubted: as Field (2009, 2015) notes this can happen with certain theories of the liar. My suspicion is that this is simply a count against such theories. But even if not, it is far from clear in these special cases whether, say, my proposed analog principles are doing any worse than Field's, precisely because of the controversies over the rational way to respond to liar-like phenomena.

[18] This is, I suspect, close to an original objection of Harman, who noted that any force of logic for reasoning comes from their being known or recognized, at which point it becomes hard to see how their import for reasoning differs substantially from any other non-logical principles.

We may again see this a little more clearly by considering the principle's implications for logical truth given by (S⁻), and a counterpart norm for truth ($\overline{\text{S}^-}$).

(S⁻) If according to S's best estimation at the time, S takes it to be the case that $\models C$ and S has reasons to consider or considers C, then S has reasons to believe C.

($\overline{\text{S}^-}$) If according to S's best estimation at the time, S takes it to be the case that C is true and S has reasons to consider or considers C, then S has reasons to believe C.

It might be objected that in this case there is an asymmetry. ($\overline{\text{S}^-}$) plainly licenses a form of bootstrapping and to that extent is implausible: it says that if I think p is true, in my best estimation, then I have reason to believe p. I am inclined to agree that ($\overline{\text{S}^-}$) is implausible on these grounds. It's just that (S⁻) is equally implausible, licensing a related form of bootstrapping: it says that if in my best estimation p is a logical truth, I have reasons to believe p. But this is intuitively untrue.

Donald believes that there is an even number of stars, but quickly realizes he has no information bearing on the question. So he reflects further: is it a logical truth that there is an even number of stars? Donald, who is horrendous at logic, thinks on this matter and (in his best—i.e. horrendous—estimation) settles the question in the affirmative. Does it follow that Donald does in fact have some reason to believe there are an even number of stars? I think not and, accordingly, that (S⁻) (and so (S)) are wrong: the fact that something holds in one's best estimation provides no reasons at all if one's best estimation is awful. Once we see this, we can see that actually some kinds of boostrapping worries afflict all of ($\overline{\text{S}^-}$), (S⁻), ($\overline{\text{S}}$), and (S) equally.

Steinberger may be happy to embrace this consequence. He bills his principle (S) as only supplying a 'directive' norm, which has "the purpose of providing first-personal guidance in the process of practical or doxastic deliberation."[19] He stresses that as such, these norms should only be held to standards consistent with the norm's serving a fruitful role in guiding reasoning. As he puts it: "It may be that the only norms sufficiently transparent to us [to be followable] are ones whose triggering conditions appeal to an agent's states or attitudes."[20] I agree that some conditions on attitudes may be appropriate for directive norms. But I think that even if we restrict our attention to directives,

---

[19] Steinberger (2019a, 316)

[20] Steinberger (2019a, 317).

Steinberger has still gone too far. Many instances of consequence and failures of consequence are not only *a priori* but blindingly obvious. Nothing prevents reasoners from holding the correct logical views on such matters beyond their logical obtuseness. To be a norm that is follow*able* in the intuitive sense doesn't mean that one has to actually take it to hold (whether explicitly or tacitly)—intuitively it need only be in reasonable epistemic reach. So as it stands Steinberger's epistemic triggering condition is simultaneously too weak (for allowing Donald to get reasons he doesn't have) and too strong (for failing to condemn equally obtuse reasoners who lack any attitudes towards obvious logical facts that bear on their reasoning).

That is my suspicion. But even if this is wrong, the more important point is the consideration of symmetry: any defense of (S^-) from bootstrapping worries is liable to save $(\overline{S^-})$—again with the latter obviously being more fundamental and general. If so, there is nothing distinctively normative about logic to be found here.

I wanted to mention one final concern for (S), which is far from a knock down consideration against it, but which will be important for my ensuing discussion. This is that principle (S) is extremely weak and qualified: it does not require conformity with logical principles, but only provides defeasible reasons for doing so (similar concerns apply to MacFarlane's (wr+)). To say that we sometimes have *some* reason to conform to logical norms is an extremely weak claim, for it is consistent with those reasons constantly being defeated. (One could, for example, claim that the goodness of golf is a reason for everyone play it. Of course, any reasons supplied would be trumped in virtually every case.)

One reason for Steinberger's retreat to reasons-based norms is that those have the virtue of giving resources to respond to the preface paradox: logical reasons for inferring a large conjunction of one's beliefs, from the individual beliefs, may be trumped by inductive reasons for humility. But accommodating some defeasibility highlights the nebulous character of the principle. Suppose I have testimony from two reliable sources about the identity of the culprit of a crime, and I believe each. But their recommendations diverge, and I have seen footage convincing me there is a single culprit. My beliefs are inconsistent—so I know from (S) that I have reason to abandon at least one of them. But: sufficient reason? Intuitively, yes. But nothing in (S) guarantees this.

Steinberger (2019a, 323–4) acknowledges this concern, suggesting that we can explain our intuitions using the idea that we have competing logical and epistemic norms, ranked by priority in ways that shift with context. In preface-paradoxes, broader epis-

temic norms outweigh logical ones, but the reverse holds in 'ordinary' cases like the one I just supplied.

Even if true, this seems more like a description of a desired solution to a problem for logic's normativity than the solution itself. Even if we get an extensionally correct theory, why are the reasons provided by logic varying in strength or efficacy? Is it that they are constant in strength, while the strength of other epistemic norms varies? Or does the strength of the reasons supplied by logic itself vary? Either way: what is that strength? All (S) tells us is that there is some. Being told that it is always enough, in interaction with other reasons, to account for our intuitions can feel dissatisfying. It is true that it is hard to find objections to such a theory. But there are real concerns that this is only because of how weak and non-committal the theory has become.

This last point may seem like an unfair objection. But I think it helps illustrate a general point about our trajectory: the program of supplying bridge principles has been marked by continual process of weakening and hedging. As we proceed from Harman through MacFarlane, Field, and Steinberger, operators are given wide scope, epistemic constraints of increasing strength are built into triggering conditions for norms, norms are weakened from strictly obligating to reason-providing. While it is true that we may be getting closer to a true principle through such hedging, there is a concern that the weak principles we arrive at are diluting or omitting something essential to logical force, which is intuitively absolute and exceptionless. We have strayed very far from that initial guiding idea.

I want to press the idea that the bridge principles we have examined have lost sight of something integral to logic's normativity with a final simple set of objections to all of them.

Consider someone who makes a series of counterfactual suppositions and then 'under supposition' affirms the consequent. They suppose q, and if p then q, then conclude under supposition: "well, therefore p." They do nothing further. This, I take it, exhibits a paradigmatically illogical form of reasoning. If there are logical norms of any kind that govern reasoning, it seems this person should have *already* violated them. They need not arrive at a belief (for example of some conditional) to have contravened logic's dictates and be pronounced a poor reasoner on specifically logical grounds. And they seem to have done something wrong, by logic's lights, that could have also occurred for belief. If someone affirms the consequent while believing, they have intuitively made a distinctively logical mistake in reasoning, and the very same one that was made by the

supposer.[21]

This raises an immediate, simple set of concerns. First, no bridge principle we have seen has any implications for supposition states. They only speak of beliefs or credences. Moreover, the principles do not seem to be extensible in any straightforward way to suppositions. Take, for example, MacFarlane's (wo-) which essentially forbids believing logical contradictions. The analog of such a principle for supposition is implausible if we are rationally permitted to suppose contradictory information for the sake of a *reductio*—a procedure which often appears highly rational. We do not need anything as recondite as the cases of PARADOX to make this point. I think similar things can be said of the other principles. (For example, is it clear that someone is *irrational* if they don't extend their suppositions logically, even when those extensions are under consideration?)

Second, even if the principles were extensible, none of them seem like they are in a position to say what has gone wrong with either the supposer, or even the believer, who has fallaciously affirmed the consequent. For example, even the believer who affirms the consequent is not necessarily condemned by any of (wo-), (wr+), (D*), or (S): the only principles from this set that forbid anything merely forbid broadly inconsistent sets of attitudes, which our reasoner never has.

To be clear on one point, the authors I have been discussing are typically not looking for norms that would govern suppositions, and are sometimes quite explicit about this. MacFarlane, for example, seems to acknowledge a possible task for logic that would have implications for mere supposition but, for reasons I will discuss critically in §4, claims that task is worth setting aside for a more theoretically fruitful focus on belief.

However, my example raises concerns for that methodology. When one person affirms the consequent under supposition, and another does it while believing, they intuitively make mistakes of the very same, distinctively logical kind. Logic impugns both believer and the supposer as bad reasoners, and on the same grounds. What this suggests is that the guiding methodological assumption should be that logic applies to the domains of the supposed and the believed equally, so that what is fundamental and distinctive of logical normativity is felt equally in both. We should of course concede that logic may have special, particular downstream effects for belief as a result of the interaction between logical norms and norms specifically governing belief (like that one should

---

[21] Authors such as HLOBIL (2015) have used examples like this to argue that there are characteristically diachronic norms of reasoning. My use of the example here is for sightly different purposes, which are compatible with taking norms to be synchronic.

believe obvious truths, or respond well to one's evidence, and so on). But a concern for all the bridge principles seen so far is precisely that they are *mixing* logical and doxastic norms, and as a result confusing what is distinctive of logical normativity. This seems all the more apparent from the fact that many clear instances of logically fallacious reasoning, even for belief, are transparently ignored by all the principles we've considered. Something has gone wrong.

3   Inferential Goodness and Bridge Principles

The discussion of §2 was obviously quick. My objections are hardly decisive: there is not only room to rebut my objections head-on, but to continue to refine the bridge principles in response to them. My goal so far has primarily been to raise suspicions—to remind that the ways in which bridge principles encounter obstacles, and are successively weakened, give us reasons to think not merely that we haven't yet found the right one, but that there is something misguided about the shape of the project as currently conceived. In this section, I want to deepen and defend that suspicion.

Note two presuppositions built into the form of every bridge principle we've encountered. First, they all involve non-evaluative terminology (*ought*, *may*, *reason*). Second, they apply this terminology to combinations of mental states (*beliefs*, *credences*). It is noteworthy that alternatives to principles of this form are hardly considered in the space of options.[22]

These presuppositions should seem especially noteworthy in light of the proposal in §1. If that proposal is on the right track, logical norms are fundamentally formulated in evaluative terms like *good* or *correct*, and the norms apply not to states, but *acts* or *processes* that mediate between them. We could, if pressed, formulate the view using a kind of bridge principle. It would look something like the following:

(Good) If $A, B \models C$, any inference from a state accepting A and B to one accepting C, in which the inference's necessary truth-preserving character is appreciated by the inferrer, is good *qua* deductive inference.

---

[22] Steinberger (2019a,b) is a rare exception in pointing out that bridge principles can be evaluated along different normative dimensions: as directives, evaluatives, or appraisals. What is intriguing is that Steinberger takes even evaluative bridge principles to be formulable with terms like *ought* or *reason*. This may be connected with the related assumption that the norms apply to agents, or their attitude states, and not acts.

Let me begin by noting two things about this principle. First, it is exceptionless, simply sidestepping all the major concerns for bridge principles including those I newly raised. Indeed, I claim the principle is more or less trivially exceptionless—it is not obvious what could count as an objection to the principle, at least provided the views of inference outlined in §1 are correct. Second, the correctness of the principle would illuminate why the counterexamples to rival principles are arising in the form that they are: the counterexamples are all the characteristic result of trying to shoehorn a fundamentally evaluative notion governing acts into (typically deontic) norms governing something like the act's performance.

To see the exceptionlessness of the principle (Good), let's quickly run through the standard concerns for bridge principles, where we will see a pattern emerge.

BACKTRACKING tells us an entailment can be reason to abandon one's starting beliefs. (Good) is silent on this issue, and so is compatible with this claim. (Good) tells us one way to infer well. It does *not* tell us when it is a good time to infer. If someone has transparently false beliefs, and applies *modus ponens* to arrive at an even more absurd conclusion, the problem is not that they have made a bad inference (that is: an inference performed badly by the standards inherent to inferring). Rather, they have made a correct inference, when the situation didn't call for an inference. They have performed an act, well, that they shouldn't have wasted their time performing. (Compare: someone can bake a cake well at a time when they should not be baking—say, they are on the verge of being consumed by a fire.)

BOOTSTRAPPING tells us that although $p \models p$, believing p provides no reasons for believing p. (Good) is silent on this issue, and so is compatible with this claim. All it says that if $p \models p$, an appreciated inference from p to p will be a good inference (which is true).

CLUTTER tells us it is irrational to clutter our minds with needless entailments of our beliefs. (Good) is silent on this issue, and so is compatible with this claim. If someone clutters up their mind by adding disjunctions to their beliefs (Good) may say that they are inferring flawlessly. It will not say that it was a good idea for them to waste their time performing those, otherwise flawless, inferences.

EXCESSIVE DEMANDS tells us that we can't fault reasoners for failing to derive far-flung consequences of their current beliefs. (Good) is silent on this issue, and so is compatible with this claim. Indeed, the far-flung consequences of our beliefs are precisely those we cannot easily appreciate. The motivating theory behind (Good) tells us we

can't correctly infer those things directly.

PARADOXES tells us that sometimes, in hard cases like the preface paradox, the liar paradox, the Sorites paradox, or cases of anti-expertise, it is rationally permissible to have beliefs not closed under simple entailment, or even to have inconsistent beliefs. (Good) is silent on this issue, and so is compatible with this claim. The form of reply here is as for BACKTRACKING. To take the preface paradox: (Good) only tells us that inferring a conjunction from its conjuncts (in an appreciable way) is to make a good inference. It does not tell us that it is good (rational, permitted, required) to make that inference.[23]

And (Good) is part of a more global view of logic and inference which handles my objections to bridge principles as well, which were based on the believer and the supposer who affirm the consequent. It even handles these cases in the same way, as I claimed would be desirable. Both the believer and supposer have performed inferences of the same type: they deductively inferred that *p* from the claims that *if p then q* and *q*, and on the basis of the claims having that form. Inferences of that form will not generally preserve truth at all possibilities, whether those inferences operate over beliefs or suppositions. Accordingly, that inference type cannot be appreciated as a necessarily truth-preserving inference on the formal grounds the reasoners have employed. As a result, the inferences of both believer and supposer are bad inferences, and bad for exactly the same reason.[24]

I said earlier that (Good) 'sidesteps' the traditional obstacles for bridge principles, and I meant it. As should now be clear, (Good) simply makes no commitments about the issues that are pressing for rival views. I also said that it is not clear what could possibly count as a counterexample to the principle, once the view of §1 is in place. The reason is simple: if inference really has as its proper function to appreciably extract information from an information state, and uses of "good" are tracking whether inference succeeds in that function, then nothing—no paradox, no odd epistemic circumstance, no quirk of reasoning—could possibly stand in the way of (Good) being true. Such

---

[23] This is of course in no way to provide any *solution* to the preface paradox. It is to say that provided there is some solution consistent with the claim that one can maintain contradictory beliefs in the preface paradox, (Good) will give logic a kind of normative force consistent with that solution.

[24] Note: it is important to condemn the inferences on their poorly appreciated formal grounds, rather than simply on the failure of the content-transitions to necessarily preserve truth. This is because the content-transitions, once we cease abstracting from the formal properties of the inference, may yet preserve truth at all possibilities (to take one example: if *p* is a lexical entailment of *q*). This is part of what makes the mistake distinctively logical.

things can only influence *when* one should exercise one's capacity to infer, not how to properly employ that capacity, once exercised.

I suspect some will see these as vices rather than virtues—indications that (Good) is too non-committal, or simply ignoring the issue we wanted to investigate. I will say more on this issue soon, especially in §4. But I want to set this concern aside briefly to highlight an important lesson. Not only does (Good) avoid the standard objections to bridge principles, but it shows that those objections are, from a certain perspective, unified. Aside from BOOTSTRAPPING, the initially disparate objection types above appear to be of one and the same *kind*. This is why the response on behalf of (Good) is essentially the same for each of those objections.

Why is it that standard objections to bridge-principles have this common form? I want to suggest that this is because the program of finding bridge principles has mistakenly been trying to shoehorn an evaluative notion governing a mental act into a deontic constraint on attitudes. We can see why this project would repeatedly encounter a single and recurring style of objection with the help of an analogy.

Consider any other evaluative notion governing an act or activity, and ask what would happen if we tried to capture it using bridge principles involving deontic language that applies to the agent's performing the activity. For example, take a fastball pitch in baseball. This is a standard form of pitch which is geared at producing a strike by testing the batter's reflexes. Given this purpose, a *good* fastball pitch is one that has (among other features) high speed and little lateral movement.

If we accept that this is what it is for a fastball to be a *good* instance of its kind, what does this tell us about what pitchers *ought to do*? It is not obvious. Consider trying to capture the evaluative notion with claims like the following.

(ia) If *A* is [known to be] a good fastball, then one ought to pitch it.

(ib) If *A* is [known to be] a good fastball, then one has reason to pitch it.

(ia) has obvious counterexamples, even restricting our attention to pitchers on the mound in a game. Perhaps the batter is fantastic at hitting fastballs, but terrible at hitting other pitches. Or, perhaps you can get a strike with your perfectly pitched fastball, but you stand to do even better—getting an out—by throwing to first where a runner is leaning too far off base.

I would be tempted to say that in some of these cases (like the first) you don't have any particular reason to pitch a fastball, falsifying (ib). But I don't need such a strong

claim for my purposes. It suffices to note that any reasons you do have to pitch a fastball are weak, defeasible, and do not capture the strength of the evaluative notion of a good fastball pitch with which we began.

We might weaken the principles further by moving from prescription to proscription.

(iia) If *A* is [known to be] a bad fastball, then one ought not pitch it.

(iib) If *A* is [known to be] a bad fastball, then one has reason not to pitch it.

Again (iia) has simple counterexamples. Maybe you are a terrible pitcher, but your (admittedly bad) fastball is the best of your bad pitches. Or maybe you're a fine pitcher of fastballs, but you could get a strike now with a poor pitch, and get the added benefit of setting up advantageously deceptive expectations for the next, better batter by doing so. In both cases, your bad fastball is the pitch you ought to make (in the second case, precisely because it is bad).

I would be tempted to say that you lack reasons to do otherwise, falsifying (iib) as well. But again, it will suffice to note that any reasons you do have to avoid pitching your fastball are weak, defeasible, and do not capture the strength of the negative evaluative notion of a bad pitch with which we began.

The pattern here should look familiar. We're seeing that when an evaluative standard governs an act, it is not easy to cash this out in terms of reasons or obligations one has to perform the act. What's more, it is easy to see precisely why this would be the case, and what kinds of counterexamples would arise for any attempt to effect that transition.

The acts we have considered as good, whether they be inferences, cake-bakings, pitches or anything else, are called "good" in connection with their associated end or purpose. It is from that end or purpose that the act derives its standard of goodness: the features of the act that promote or secure that end. From this two things follow.

First, the standard of goodness being applied is tracking features relative to that fixed end. As a result, calling the act "good", in this sense, has no implications for whether the act should be performed or not. Saying that a pitch is a good one is not to say you should pitch it. That depends on whether the situation calls for that kind of pitch. Likewise, saying that an inference is good (*qua* inference), or performed well, or correctly, is not to say that one should perform it. That depends on whether the

circumstance calls for an act of information extraction. There is simply no tension between saying that such-and-such is what it takes to perform an act-type well, but that the circumstances don't call for that act-type at present.

Second, there is a simple recipe for finding counterexamples to the claim that when an act is good as its act type one should perform it: find reasons against promoting or accomplishing the act's goodness fixing end. Do you want to find a case where, even though a pitch of a fastball is good *qua* fastball pitch, you shouldn't pitch it? Easy: find a case where you have no reason to try to test a batter's reflexes by getting a strike. This can be because the batter's reflexes are too good, or you have better things to do (like getting an out by throwing to first). Or, more simply, it could be because you reasonably can't perform the pitch well at all.

Do you want to find a case where, even though an inference is good, *qua* inference, you shouldn't execute it? Easy: find a case where you have no reason to perform a total act of information extraction on an acceptance state. That could be because one is in a position to see that basing acceptance state contains false information, and the concluding acceptance state is somehow regulated by truth (Backtracking); or because even though the information in the acceptance state seems reliable, you have more useful things to do (Clutter); or because you're in a tough epistemic situation where extracting the information from your belief state is going to lead to foreseeably incorrect, and otherwise pernicious, acceptance states (Paradox). Or maybe more simply, just find a case where you are (reasonably) not in a position to make the inference well due to your own limitations (Excessive Demands).

So here is the final lesson: if the normativity of logic is as described in §1, we can see that the counterexamples to bridge principles are almost all of a single, predictable form. They are all the very sort of objections one would encounter if one were mistakenly try to take an evaluative normative notion governing an act, and transpose it to illegitimately draw conclusions about when one ought, or ought not, perform the act.

One could, of course, adjust the resulting norms (whether they concern inference, or pitching, or cake-baking, or any other act) to avoid the counterexamples. But one could only do this at the expense of losing sight of the original norm governing the goodness of the act. The way to do this would be to start encrusting distinct conditions that track the norms governing not, or not only, the goodness of the act, but the conditions which make it reasonable to perform the act. In the case of inference, at least if we fixate (illegitimately) on cases involving only beliefs, these encrusted conditions will

start to tack on features that make beliefs reasonable in light of their truth, or apparent truth. And, as I've already noted, this is precisely what we seem to see in the progression of principles building on Harman's starting point.[25]

As such, the importance of the view of §1 is not merely that it arrives at a conception of the normativity of logic which avoids the existing barrage of counterexamples to bridge principles. Nor is it even that the view seems to stand immune from any similar form of counterexample. Rather, the view illuminates key methodological presuppositions—presuppositions that can and should be questioned—that seem to lie at the heart of the discontents of existing treatments of logic's normativity.

4   Logic and Reasoning

In §3, I set aside two related concerns that it is now time to take up. First there is a concern that my view of logic's normativity is weak, precisely because it avoids the threat of standard counterexamples by being non-committal. Second, there is the related concern that my view is not properly engaging with any of the concerns about logic that lead philosophers like Harman, MacFarlane, Field, or Steinberger to investigate logic's normativity.

This second concern is especially pressing, since it is unclear whether the authors I engage with even have a common theoretical goal in providing bridge principles.[26] Fortunately, despite the differences between their approaches, there is at least one unifying thread among most discussions of bridge principles: a concern with logic's relevance to *reasoning*. Accordingly, the first step in assessing how the view of §1 is engaging in the dialectic is to say what implications that view has for reasoning.

We should begin with Harman. Harman, I noted, supplied bridge principles mainly to emphasize their susceptibility to counterexample. As Harman (1986, 5) puts it: a "logical principle [i.e., a logical relation of consequence or what Harman calls an "implication"] holds without exception, whereas there would be exceptions to the corresponding principle of belief revision"—where principles of belief revision are en-

---

[25] Indeed, in trying to formulate principles explaining merely how truth governs belief, one finds precisely the same kinds of moves as in the normativity of logic literature: wide-scoping, strengthened epistemic triggering conditions, and various contortions to avoid paradox. One can in fact see all of these moves made in the survey of truth-governed norms for belief in Bykvist & Hattiangadi (2007).

[26] See in this regard the discussion of Steinberger (2019b).

shrined in bridge principles.[27] For Harman, this contributes to a central contention that logic is "not of special relevance" to a theory of reasoning, where "reasoning" is interpreted broadly to involve general procedures for revising one's beliefs, including abandoning them.[28]

Why did Harman take the defeasibility of bridge principles to contribute to the claim that logic is not of special relevance to reasoning? Harman certainly does not think that logical rules of implication have no importance for reasoning at all. Indeed, he takes them to stand as integral truths that help regulate what we should and should not believe, albeit defeasibly. But as such it is not clear what makes logic more relevant to reasoning than other suitably general and stable truths. For example, as Harman notes in regard to the issue of inconsistency, "[p]rima facie, one should not continue to believe things one knows cannot all be true, whether this impossibility is logical, physical, chemical, mathematical, or geological."[29] This is not Harman's only case against logic's special relevance to reasoning, but it is central.

I may seem to have been siding with Harman on the issue of reasoning, perhaps emphasizing a different role for logic to play. After all, a central idea of my discussion is that logic helps track conditions when an inference is performed well or correctly, but that doing this should be sharply separated from any claims about *when* to perform an inference—even a good one. It is clear that Harman is keenly interested in the latter kind of question.

The claim that I side with Harman is partially right, but it is worth noting that the dialectical situation is somewhat complex. Let me begin with some ways in which I agree with Harman. First, I've conceded that logic doesn't give information about how to reason, in the sense of when to engage in reasoning of certain kinds. Instead logic has implications for how to correctly perform an act that is part of reasoning— namely deductively inferring. What's more, in investigating conditions on good deductive inference, deductive logic only investigates one of many such activities of reasoning (including inductive inference), it tracks the goodness of this activity imperfectly (by ignoring lexical entailments), and it only investigates a necessary but insufficient condition on that goodness (by setting aside appreciation). Harman himself emphasized

---

[27] Harman sometimes calls these latter principles "rules of inference" (HARMAN, 1984, 108). It should be borne in mind that his "inference" is tracking reasoning broadly construed, not the process I have called "inference" in §1.

[28] HARMAN (1986, 11).

[29] HARMAN (1984, 109). See also the discussion at HARMAN (1986, 17).

points similar to all these three in attacking logic's special relevance to reasoning.

But when Harman expands on his claim that logic lacks special relevance for reasoning, he sometimes goes too far. Inference, as just noted, is *part* of reasoning. It is a central, if not essential, such part. A reasoner simply could not get by, as a reasoner, in any ordinary course of existence without the ability to draw deductive inferences. (Imagine an agent seeking food who knows that if the prey didn't go down path A it went down path B, and that it didn't go down path B, but is unavoidably stuck in attitudinal limbo, not merely because they can't see the goodness of the inference, but because inferring isn't even in their mental repertoire.) I think it is even an open question whether, if an agent lacked the ability to draw deductive inferences, we could even consider them a reasoner, or a thinker more broadly.

It is thus highly misleading to portray logic's relevance to reasoning as being like that of, say, physics, chemistry, or geology. It is highly misleading to compare the implications of logic to those like "*X plays defensive tackle for the Philadelphia Eagles* implies *X weighs more than 150 pounds.*"[30] Physics, chemistry, and geology, for instance, provide truths that one *reasons with* (as arguably does the inductively supported general truth about Eagles defensive tackles). Such truths do not, properly speaking, constrain reasoning processes, but merely furnish the materials for reasoning with. The standards of goodness governing inference normatively constrain a process of reasoning itself. An agent with no knowledge of the physical, chemical, or geological sciences, or of American football, can still be an excellent, indeed perfect reasoner. But being unable to draw deductive inferences well would be devastating to any reasoner—possibly even precluding them from counting as a reasoner at all.[31]

To clarify this point, it may be helpful to develop another analogy. Suppose some-

---

[30] HARMAN (1986, 17).

[31] Part of what is holding Harman back, at least in his early writing, is that he finds himself unable to rule out a view on which logic merely consists of a body of truths, distinguished at most by their generality. (Harman is clear, at least in HARMAN (1984) that he also finds himself unable to completely agree with such a view.) He considers against this idea only an argument from Carrollian regress. I am not even sure we need arguments against this view in the current dialectic—I think it is perfectly reasonable to take the rival view that logic tracks necessary-truth-preserving entailment relations as a starting point barring further argument. But if we need argument, we can try to do so from the ground up as in §1. Alternatively, we can simply attack the rival construal of logic on its own terms—see in this regard especially ETCHEMENDY (1990) who (rightly or wrongly) attributes a view like that Harman discusses to Tarski. Another obstacle is that Harman only engages with the 'acceptance' of logical rules either in terms of belief, or in terms of brute dispositions. But there are noteworthy alternatives (see especially the above literature cited in nn.6,7).

one said: "Being able to bluff well is of no *special* relevance to playing poker well." I suspect poker aficionados would take exception to such a claim, at the very least in its wording. But we can concede that there is one understanding of this claim which is true: it's not as if bluffing is all one does in poker. To play poker, and especially to play it well, one often simply has to openly play the strength of one's current hand, or even fold.

But suppose our character continued: "What's more, being able to bluff well is of no more relevance to playing poker well than having a good hand." At this point, our speaker has gone beyond a potentially misleading statement into confusion. A poker hand is what one plays poker *with*. One can play well with a bad hand, or play poorly with a good one. To lump the possession of a good poker hand in with the skilled actions of strategically betting or folding shows a serious confusion about the nature of the target of investigation.

In claiming that logic is no more relevant to good reasoning than the sciences, Harman has not merely understated the importance of logic to reasoning, but confused the distinctive way in which logic contributes to that study. Truths are what one reasons with, including through deductive inference, just as a hand is what one plays poker with, including by bluffing. One can reason poorly with truths, and well with untruths, just as one can play poorly with a good poker hand, and well with a bad one. Harman's claims about logic's lack of special relevance to reasoning are sometimes founded on a conflation of the activities of reasoning with their objects. If we think in this way, we are apt to miss the one very way in which logic should be viewed as special for the study of reasoning. As Etchemendy aptly put it, in a different but related context: "Logic is not the study of a body of trivial truths; it is the study of the relation that makes deductive reasoning possible."[32]

So although I agree with Harman on many points, I think he sometimes goes too far in trying to downplay the importance of logic in the study of reasoning. Logic has what I would think of as a very significant role in that investigation: it studies a huge class of content-transitions that undergird a necessary condition on performing a central activity of reasoning—deductively inferring—well. It studies, (indirectly, and with certain limitations) how to correctly perform an action that is partially constitutive of reasoning well. Whether this is a 'special' role is a vague matter. But one cannot, as Harman occasionally does, compare this role to knowledge of certain truths, even important and

---

[32] ETCHEMENDY (1990, 11).

general ones, which is no part of good reasoning to begin with. The problem with Harman's discussion of logic's role in reasoning was not, as many seem to suppose, that he failed to adequately refine his considered bridge principles, but that he failed to properly locate the distinctive *kind* of contribution that logic makes to the study of reasoning—a contribution which isn't formulable in terms of principles constraining combinations of attitudes.[33]

What of other authors? MacFarlane's discussion is perhaps the most instructive to consider. For MacFarlane seems to recognize something very close to the normative role for logic that I have set out. But as soon as he notes it, he sets it aside as straightforward and unilluminating.

MacFarlane cites an interest in getting clear on the normative role of logic in reasoning as opening up a way to arbitrate logical disputes, both over choice of logic and over foundational questions in logic. Like Harman, he is careful to distinguish differ-

---

[33] I see similar issues for a contention of RUSSELL (2017, §4) and BLAKE-TURNER & RUSSELL (forthcoming, §3), who claim that logic is not normative in any interesting sense because core logical statements are descriptive, and only have normative consequences alongside other normative assumptions.

> In conjunction with common normative commitments concerning truth and falsity (only believe what is true, don't reason to false conclusions, etc.), logics …have normative consequences …if this is how logic is entangled with the normative, then it shares this status with paradigmatically descriptive scientific theories, including those of physics and mathematics. (RUSSELL, 2017, 10-11)

I agree that if the normative import of logic only comes via norms governing truths, there would be nothing that would set it apart from physics or mathematics. (Indeed, I've argued against several theorists like Field and Steinberger above making essentially that point.) But the problem Russell and Blake-Turner rightly identify is not merely arising (as I think they intimate) from the fact that normative consequences for logic only result from pairing its descriptive claims with some separate normative commitments or other. It is rather the presumption that *normative commitments about truth in particular must be doing the work*. That is what would lump logic in with mathematics and physics. There are other normative commitments, besides those involving truth, that bring out constrained subsets of descriptive truths as having *distinctive* normative import. Suppose, for example, that act utilitarianism is true. Then the fact that such-and-such act is utility-maximizing may be a purely descriptive one, that only has normative consequences in conjunction with act-utilitarian principles. But it would be extremely misleading to say in this context that facts about utility-maximization were normative in a way no different than those of mathematics or physics. On the view I'm putting forward, the primary point of investigating descriptive facts about content through logical theory is that these have distinctive import for the goodness of a mental activity. This import is not shared by the truths of physics and mathematics, and it would accordingly be confused to treat logic's normative status as on a par with theirs.

ent things one could mean by "inference" or "reasoning". And in making one such distinction he says the following:

> In a more formal sense, reasoning is a process of drawing out the conse-
> quences of a given set of premises. One need not believe the premises: one
> might just be investigating them, or using them in a conditional proof or
> *reductio ad absurdum*. To distinguish this process from reasoning in the
> sense of "reasoned change in view," we might call it "inferring" (though
> "inferring" may be subject to the same kind of ambiguity as "reasoning").
>   …I think it is relatively uncontroversial that logic provides norms for
> inferring (in the narrow sense of drawing out consequences). For the proof
> rules of a logic are *explicitly* normative: for example, the ⊃-elimination rule
> says that if you have already written down *A* and *A* ⊃ *B*, you may write
> down *B*. These proof rules license or permit certain inferences.
>   …So here is a clear sense in which logic is normative for reasoning. But
> this sense isn't going to help us much with the problems we looked at in
> the last section. Our intuitions about when it is permissible to infer a con-
> clusion from some premises (in the narrow sense) have the same sources
> as our intuitions about logical validity: primarily, our logical training. (In-
> deed, it takes some logical training in order to engage in the practice of
> "inferring" at all: one must be trained not to use information not con-
> tained in the premises, for instance, and not to worry about whether the
> premises are true.) Thus these intuitions are likely to be subject to just the
> same "indoctrination biases" as our intuitions about validity. A classicist
> will take it to be correct to infer anything from a contradiction in formal
> argumentation, while a relevantist will not. If we are to get beyond this
> kind of conflict of intuitions, we need to talk about norms for reasoning
> in the broader sense: norms for belief and belief change.[34]

To evaluate MacFarlane's claims here, I need to draw an added distinction: that between formal reasoning through the use of a particular deductive system of the sort that is taught in a logic class (perhaps on the added assumption that it is a 'correct' one) and the mental activity I discussed in §1. I reserve the term "inference" for the latter, and will call the former "symbolic reasoning". I think MacFarlane could have meant either

---

[34] MacFarlane (ms/2004).

of these two things by his "inference". Accordingly, it is worth considering how his remarks fare on each interpretation (without saddling him with either one).

Much of what MacFarlane says in this passage holds true of symbolic reasoning. For example, the claim that proof rules are normative in the sense that they preclude certain ways a deduction may proceed, but allow others. Also it seems true that it takes training to engage in the process, and that this training may open up theorists to biases toward certain proof theoretic frameworks. I take issue with none of these claims.

But if we were to interpret MacFarlane as talking about what I call "inference", some of the corresponding claims would be true, while others are problematic. MacFarlane describes the target phenomenon as a "process of drawing out the consequences of a given set of premises" which is very close to my characterization of inference as an information-extracting act. And as I've emphasized, information-extraction does not necessarily operate on belief states, so that it is not necessary that one believe anything while inferring—again according with MacFarlane's characterization.

But, importantly, the claims of the last paragraph quoted above would be untrue. It is not true that it "takes some logical training in order to engage in the practice of "inferring" at all." This claim is not clearly true even if we restrict our attention to deductive inference. Such inferences have been made continuously by mathematicians throughout history, not all of whom had particular training in logic. It is also possible such inferences are made by ordinary persons regularly, in simple applications of rules like *modus ponens*. Indeed, I think it is an open question whether animals engage in inferences, and even deductive ones.

What is true is that one needs instruction to *theorize about* good inference. In this respect, processes of inference might loosely resemble, say, the kinds of processing that go on in composing and parsing syntactic structure. This composing and parsing is something we do all the time. And, we can suppose, we often do it correctly. But it may still take hard theoretical work to say precisely what it is to do it correctly.

Are our intuitions about what makes a deductive inference good subject to indoctrination biases? Probably. And in this respect perhaps one could maintain there is much more bias in the logical or broader inferential case than in the case of syntax. But one thing that MacFarlane is stressing is that we somehow improve our situation by discussing broader norms for reasoning. But this suggestion, if made in respect to the study of logic's normativity within the sphere of what *I* have been calling "inference", is as far as I can see not only unhelpful, but actually counterproductive.

The suggestion is unhelpful because, as stressed in my discussion of Harman, inference is a constitutive and central component of reasoning more broadly construed. In reasoning, and in order to reason, one must sometimes infer.[35] What this means is that even if we could pinpoint the norms for reasoning more broadly construed, one of two things would be true. Either those norms wouldn't happen reflect the influence of the norms properly governing inference, or they would. If the former, then the norms actually end up telling us nothing pertinent to logic—they concern those aspects of reasoning which have no bearing on logical domains. If the latter, then those norms are bound to be just as controversial, if not more controversial, than the simple normative claims about inference itself. There is no reason to think that when the goodness of a particular belief-forming strategy critically turns on whether a particular inference is a good one, that the belief-forming strategy is going to be any less controversial, or influenced by indoctrination biases, than the inference considered in isolation.[36]

Indeed, this thought not only reveals that MacFarlane's proposal to consider norms of reasoned change in attitude is unhelpful, but in fact counterproductive. Once we move to investigations of reasoning more broadly construed, we will be investigating a mixture of logical and non-logical norms. If we focus especially on norms governing reasoning with full beliefs, for example, those norms will certainly integrate broader norms for belief formation, that may have nothing to do with inference in particular, and so nothing to do with logic. This is precisely what I argued has happened to investigations of the normativity of logic in §2: traditional bridge principles, to the extent they avoid counterexamples, start to encrust norms that concern things like proper responsiveness to truths. Precisely how those norms of truth-responsiveness should be formulated, especially in puzzle cases like those in Paradox, is itself a tricky matter. So there is interesting work to be done here. But the more of it we do, the more we obscure the distinctive role of logic.

As I say, it is an open question whether MacFarlane was merely concerned with what

---

[35] This is even true if we posit, as I do not think is true, that inference only operates suppositionally. For the norms governing the inference under supposition will influence how that suppositional reasoning eventually has downstream consequences for belief.

[36] One of the applications of Steinberger (2016) is to use an investigation of bridge principles to diffuse relevantist attacks on *ex falso*. This application can persist, but on different grounds. The problem is that relevantists themselves are illegitimately, sometimes implicitly, invoking bridge principles in their arguments that are not only subject to counterexample, but do not clearly bear on logic proper. The right way to object to relevantists is not to look for better bridge principles, but to note that their use of them to motivate logical restrictions is illegitimate.

I have called "symbolic reasoning", or whether he had an interest in the phenomenon I have been calling "inference" (or yet something else). Either way, his motivation for looking at reasoning broadly construed is flawed. Reasoning broadly construed comprises the phenomenon that logic studies as an integral, but proper, part. Trying to find general norms for reasoned belief revision not only fails to avoid controversies proper to the logical domain, but only obscures the normativity of logic by mixing its controversies with those concerning belief formation that need have nothing to do with logic in particular. Accordingly, it is unclear what insight into logic could be gained through the suggested strategy.

Similar worries can be raised for the investigations of Steinberger and Field, sometimes in ways that are exacerbated by the concerns of §2. Steinberger bills himself as following Harman in exploring the normative import of logic for reasoning in general, with a focus on how such norms could figure as *directives* which, recall, "have the purpose of providing first-personal guidance in the process of practical or doxastic deliberation."

I argued that the norms Steinberger ends up with are frustratingly weak, positing only the existence of some logical reasons, in some very special circumstances, of largely unspecified strength. But the important thing to note is that once we recognize that inference is a component process of reasoning, it is with respect to inference, and not belief, that we should first look for guiding principles that distinctively owe their force to logic. I do not have the space to explore such principles in any detail here, but it may be worth noting the starting point. Some obvious candidates for directives are: "Never infer badly" or "If one ought to infer, make a good inference." Logic, of course, only provides an indirect and limited investigation of the goodness or badness at issue. Still, such principles do not succumb to the problems from Backtracking, Bootstrapping, Clutter, Excessive Demands, or Paradoxes for the same reasons as (Good). And they also extend the directive normativity of logic to the merely supposed, in the same way as (Good).

Finally, consider Field. Field (2015) bills his credal norm as contributing to a "conceptual role" for the term "valid", thereby giving us insight into a kind of common core of the concept that is what is up for dispute in logical debate. As he puts it: "*disagreement about validity (insofar as it isn't merely verbal) is a disagreement about what constraints to impose on one's belief system.*"[37]

---

[37] Field (2015, 42), Field's emphasis.

For reasons noted in §2, this is dubious. The constraint given by Field doesn't seem to have any special connection with logic. Recall what his principle tells us about logical truth, and how this compares to plausible constraints on simple truth:

(D^-) If it's obvious that $\models B$, then one ought to impose the constraint that $P(B)$ is 1.

($\overline{D^-}$) If it's obvious that B is true, then one ought to impose the constraint that $P(B)$ is 1.

Certainly two people can agree about what the obvious truths are, but disagree about which of them are logical. Surely they will agree that they should have high credence in all the obvious truths, regardless of whether they are logical. So I find it hard to see what illumination Field's conceptual role can provide for the existence of a distinctively logical form of disagreement here.

There is more to say, but let me leave off discussion in order to return to the opening question of this section: Is my proposal engaging with those of my targets? What I've tried to argue is that the situation is complex. In one important sense, I am not engaging directly with the concerns of the authors I have discussed. Those authors are all concerned with general norms for reasoned change in belief states. And I am, by my own admission, providing no such general norms. This is effectively how the view I give is able to effortlessly sidestep the slew of traditional counterexamples to bridge principles.

But in another more important sense I take myself to be engaging with my targets quite directly. These authors are looking for norms governing reasoning on various motivating grounds: to gain insight into the special importance of logic for good reasoning, including its directive normative force; to arbitrate logical disputes; and to find a core concept underlying non-verbal logical dispute. In each case, I've argued, the focus on finding norms for reasoning broadly construed is a mistake. And the source of the mistake partly owes to inattention both to the plausible role logic has to play in studying inference, and to the distinctive role inference itself plays in reasoning. So the importance of studying logic through the lens of inference is not merely to clarify and circumvent the obstacles encountered on current prevalent methodologies, but to reveal how those methodologies are not conducive to accomplishing the aims of the theorists that adopted them.

5   Taking Stock

I've claimed that logic's normative import is most fundamentally directed at evaluating inferential acts and that, once this is seen, the normative force in question is simple and exceptionless. No element of this view is new. Logic texts often cash out the normative implications of validity in evaluative terms, and we have an extensive literature explaining the function of evaluative talk in precisely the ways I recommend. What is more, the default view, including the view of most authors I've discussed, is that logic is relevant to reasoning which is itself a process, or activity, or event, and not a state or group of states. And there is of course an extensive and growing literature on inference itself, in which inference is customarily viewed as just such an activity, process, or event.

So why have investigations of logic's normativity seemingly ignored this perspective? My suspicion is that the trajectory of the literature owes to a single, simple, and understandable move of Harman's. In asking how logic could be relevant to reasoning, Harman formulated his principles in terms of constraints on attitude states. This was a reasonable thing to do, given that Harman was partly moved by the attractions of a view of logic on which it merely supplies us with general truths, and a view of inference on which there was very little logically distinctive of it as a process (as noted in n.31).

I think it was certainly reasonable to explore this as one avenue among many. But I've also claimed it was ultimately a mistake. Logic doesn't merely supply us with general truths, but necessary truth-preserving relations among contents. And it does this precisely because this is a necessary condition on good deductive inference. To the extent logic has distinctive normative implications for reasoning, they apply at the level of inferential transitions.

But once Harman's reasonable move was made, it kicked off an equally reasonable research program: that of finding bridge principles capturing logic's influence on acceptable combinations of beliefs or credal states. This program is not merely reasonable because of the reasonableness of Harman's starting point. Rather, it is because there are no reasons I know of (including no reason from anything I have said in this paper) to suspect that this program cannot reach an end. There is no reason to think that we cannot formulate a bridge principle of the sort that MacFarlane catalogues which, encrusted with enough caveats and conditions, is true. Indeed, I see the current literature as making steady progress toward that very goal.

So it is important to note that my objection to this program is not that it is somehow unfulfillable. It is rather that it has nothing to do with what is distinctive of logic and

its normative relevance to reasoning. Instead, the program has much more to do with a great admixture of epistemic norms, including especially norms governing reasonable formation and management of beliefs in response to evidence and recognized truths, to which logic may contribute partially and indirectly (mostly by supplying us with certain truths, or organizing our evidence). One thing this means is that finding the true bridge principles will probably be a convoluted and frustrating task, with the resulting principles becoming increasingly qualified, precisely as we have seen. But more importantly, even if this task reaches an end, there will be no way to 'factor out' of our final principles what was distinctively contributed by logic. By the end, logic's role will be completely swamped by the complementary, overlapping, and interacting non-logical epistemic norms.

This concern is most easily seen from the fact that standard bridge principles, out of the starting gate, seemingly give up on saying why the resources of logic condemn logically fallacious reasoning. As we saw, no bridge principle seems to come close to explaining what is wrong with affirming the consequent. By now it should be clear why this is so. What is wrong with the person who has engaged in such a fallacy has nothing to do with their combination of attitudes, even over time. What is wrong with them is the process or event that led them to those attitudes. That simple idea was understandably set aside with Harman. What I am recommending is that any serious investigation of logic's normativity must take it up again.

References

Blake-Turner, Christopher & Gillian Russell. forthcoming. "Logical Pluralism Without the Normativity." *Synthese*. doi:10.1007/s11229-018-01939-3. [29]

Boghossian, Paul. 2014. "What is Inference?" *Philosophical Studies*, vol. 169 (1): 1–18. [4]

Bolinger, Dwight. 1968. "Post-posed main phrases: An English rule for the Romance Subjunctive." *Canadian Journal of Linguistics*, vol. 14: 3–30. [2]

Broome, John. 1999. "Normative Requirements." *Ratio*, vol. 12 (4): 398–419. [9]

—. 2014. "Comments on Boghossian." *Philosophical Studies*, vol. 169 (1): 19f125. [4]

Buckareff, Andrei A. 2005. "How (Not) to Think About Mental Action." *Philosophical Explorations*, vol. 8 (1): 83–89. [3]

Bykvist, Krister & Anandi Hattiangadi. 2007. "Does Thought Imply Ought?" *Analysis*, vol. 67 (296): 277–285. [25]

CAIE, MICHAEL. 2013. "Rational Probabilistic Incoherence." *Philosophical Review*, vol. 122 (4): 527–575. [12]

CARROLL, LEWIS. 1895. "What the Tortoise Said to Achilles." *Mind*, vol. 4 (14): 278–280. [6]

ETCHEMENDY, JOHN. 1990. *The Concept of Logical Consequence*. Harvard University Press. [27], [28]

—. 2008. "Reflections on Consequence." In *New Essays on Tarski and Philosophy*, DOUGLAS PATTERSON, editor, 263–299. Oxford University Press. [7]

FARKAS, DONKA. 1985. *Intensional Descriptions and the Romance Subjunctive*. Garland, New York. [2]

FIELD, HARTRY. 2009. "What is the Normative Role of Logic?" *Aristotelian Society Supplementary Volume*, vol. 83 (1): 251–268. [4], [9], [14]

—. 2015. "What Is Logical Validity?" In *Foundations of Logical Consequence*, COLIN R. CARET & OLE T. HJORTLAND, editors, 33–70. Oxford University Press. [14], [33]

FINLAY, STEPHEN. 2004. "The Conversational Practicality of Value Judgment." *The Journal of Ethics*, vol. 8: 205–223. [3]

—. 2014. *A Confusion of Tongues*. Oxford University Press, New York. [3]

FREGE, GOTTLOB. 1879?/1979. "Logic." In *Posthumous Writings*, 1–8. Basil Blackwell, Oxford. [4]

FRIEDMAN, JANE. 2013. "Question-Directed Attitudes." *Philosophical Perspectives*, vol. 27 (1): 145–174. [2]

GIBBONS, JOHN. 2009. "Reason in Action." In *Mental Actions*, LUCY O'BRIEN & MATTHEW SOTERIOU, editors, 72–94. Oxford University Press. [3]

GLANZBERG, MICHAEL. 2015. "Logical Consequence and Natural Language." In *Foundations of Logical Consequence*, COLIN CARET & OLE HJORTLAND, editors, 71–120. Oxford University Press. [7]

GROENENDIJK, JEROME & MARTIN STOKHOF. 1997. "Questions." In *Handbook of Logic and Language*, J. VAN BENTHAM & A. TER MEULEN, editors, 1055–1124. Elsevier Science, Amsterdam. [2]

HARMAN, GILBERT. 1984. "Logic and Reasoning." *Synthese*, vol. 60 (1): 107–127. [8], [9], [26], [27]

—. 1986. *Change in View*. MIT Press. [8], [9], [25], [26], [27]

Heim, Irene. 1992. "Presupposition Projection and the Semantics of Attitude Verbs." *Journal of Semantics*, vol. 9 (3): 183–221. [2]

Hieronymi, Pamela. 2009. "Two Kinds of Agency." In *Mental Action*, Lucy O'Brien & Matthew Soteriou, editors, 138–162. Oxford University Press. [3]

Hlobil, Ulf. 2014. "Against Boghossian, Wright and Broome on Inference." *Philosophical Studies*, vol. 167 (2): 419–429. [4]

—. 2015. "There Are Diachronic Norms of Rationality." *Thought: A Journal of Philosophy*, vol. 4 (1): 38–45. doi:10.1002/tht3.155. [18]

—. 2019. "Inferring by Attaching Force." *Australasian Journal of Philosophy*, vol. 97 (4): 701–714. doi:10.1080/00048402.2018.1564060. [4]

Locke, John. 1690/1979. *An Essay Concerning Human Understanding*. Oxford University Press. [4]

MacFarlane, John. ms/2004. "In what sense (if any) is logic normative for thought?" Unpublished. Delivered at the American Philosophical Association Central Division meeting. [4], [8], [9], [30]

Mele, Alfred. 2009. "Mental Action: A Case Study." In *Mental Actions*, Lucy O'Brien & Matthew Soteriou, editors, 17–39. Oxford University Press. [3]

Millikan, Ruth G. 1984. *Language, Thought and Other Biological Categories*. MIT Press. [4]

Peacocke, Christopher. 2008. "Mental Action and Self-Awareness." In *Mental Action*, Lucy F. O'Brien & Matthew Soteriou, editors, 358–376. Oxford University Press. [3]

Russell, Bertrand. 1920/1988. "The Nature of Inference." In *The Collected Papers of Bertrand Russell, vol.9: Essays on Language, Mind, and Matter, 1919–26*, J. G. Slater & B. Frohmann, editors. Unwin Hyman, London. [4]

Russell, Gillian. 2017. "Logic Isn't Normative." *Inquiry: An Interdisciplinary Journal of Philosophy*, 1–18. doi:10.1080/0020174x.2017.1372305. [29]

Sainsbury, Roy M. 2002. "What Logic Should We Think With?" In *Logic, Thought and Language*, Anthony OHear, editor, 1–17. Cambridge University Press, Cambridge. [4]

Sánchez-Miguel, Manuel García-Carpintero. 1992. "The Grounds for the Model-Theoretic Account of the Logical Properties." *Notre Dame Journal of Formal Logic*, vol. 34 (1): 107–131. [7]

Setiya, Kieran. 2013. "Epistemic Agency: Some Doubts." *Philosophical Issues*, vol. 23 (1): 179–198. [3]

STALNAKER, ROBERT. 1984. *Inquiry*. Bradford Books, MIT Press. [2]

STEINBERGER, FLORIAN. 2016. "Explosion and the Normativity of Logic." *Mind*, vol. 125 (498): 385–419. [32]

—. 2019a. "Consequence and Normative Guidance." *Philosophy and Phenomenological Research*, vol. 98 (2): 306–328. doi:10.1111/phpr.12434. [9], [15], [16], [19]

—. 2019b. "Three Ways in Which Logic Might Be Normative." *Journal of Philosophy*, vol. 116 (1): 5–31. doi:10.5840/jphil2019116n. [19], [25]

STRAWSON, GALEN. 2003. "Mental Ballistics or the Involuntariness of Spontaniety." *Proceedings of the Aristotelian Society*, vol. 103 (3): 227–257. [3]

STROUD, BARRY. 1979. "Inference, Belief, and Understanding." *Mind*, vol. 88 (350): 179–196. [4]

THOMSON, JUDITH JARVIS. 1965. "Reasons and Reasoning." In *Philosophy in America*, MAX BLACK, editor, 282–303. Cornell University Press, Ithaca, N.Y. [4]

—. 2008. *Normativity*. Open Court Press, Chicago. [3]

WRIGHT, CRISPIN. 2014. "Comment on Paul Boghossian, "What is Inference"." *Philosophical Studies*, vol. 169 (1): 27fi37. [4]

ZIFF, PAUL. 1960. *Semantic Analysis*. Cornell University Press, Ithaca. [3]