Epistemic Characterizations of Validity and Level-Bridging Principles*

Joshua Schechter

ABSTRACT

How should we understand validity? A standard way to characterize validity is in terms of the preservation of truth (or truth in a model). But there are several problems facing such characterizations. An alternative approach is to characterize validity epistemically, for instance in terms of the preservation of an epistemic status. In this paper, I raise a problem for such views. First, I argue that if the relevant epistemic status is factive, such as *being in a position to know* or *having conclusive evidence for*, then the account runs into trouble if we endorse certain familiar logical principles. Second, I argue that if the relevant epistemic status is non-factive, such as *is rationally committed to* or *has justification for believing*, then a similar problem arises if we endorse the logical principles as well as a sufficiently strong epistemic "level-bridging" principle. Finally, I argue that an analogous problem arises for the most natural characterization of validity in terms of rational credence.

1. Introduction

On a natural picture of inference, there are broadly two kinds of inference. Deductive inference is the kind of inference in which it is purported that the premises conclusively support the conclusion. Inductive inference is the kind of inference in which it is purported that the premises support the conclusion, but not conclusively. A deductive inference counts as valid if the premises do conclusively support the conclusion. It counts as invalid otherwise.[1]

Valid inferences include logically valid inferences, such as the inference from the claims that it's raining outside and that if it's raining outside then the streets are wet to the claim that the streets are wet. Valid inferences also include what might be called "analytic" or "conceptual"

* Forthcoming in *Philosophical Studies*.

[1] On some views, including those inspired by Harman's (1986) emphasis on the distinction between inference and implication, validity concerns an abstract relation among propositions, not any kind of reasoning. So it is a mistake to think of inferences as valid or invalid. If that is right, then the issues raised here should be put in terms of whether a purported entailment is valid or invalid. This would only require minor changes in the discussion to follow.

inferences, such as the inference from the claim that Riley is a rooster to the claim that Riley is a male chicken, and the inference from the claim that Sherlock knows that Moriarty committed the dastardly deed to the claim that Moriarty did in fact commit the dastardly deed.[2]

It is standard to define validity using the notion of truth, or perhaps the notion of truth in a model. For instance, one might define an inference to be valid if there is a guarantee (of the right sort) that the inference is truth-preserving. Or one might define an inference to be valid if it preserves truth in all models that conform to certain constraints (e.g., they must respect analytic or conceptual truths). But these are not the only options for defining validity. One might instead define validity in epistemic terms. For instance, one might define an inference to be valid if there is a guarantee (of the right sort) that the inference preserves knowledge. Or one might define an inference to be valid if there is a guarantee that it preserves rational commitment, or justified belief, or rational credence, or something of the sort.

In fact, several philosophers have presented definitions of validity (or a closely related notion) in terms of the preservation of an epistemic status. For instance, Sundholm (2012, p. 950) defines validity in terms of the preservation of knowledge. Prawitz (2015, p. 73) defines "legitimate inference" as the kind of inference that preserves the property *having conclusive evidence for*. Brandom (1994, p. 168; 2008, p. 120) defines "committive inference" as the kind of inference that preserves the property *being committed to*.

The aim of this paper is to raise a difficulty for the project of defining – or, more generally, characterizing – validity in epistemic terms. I will focus on proposed accounts of validity according to which an inference is valid just in case it preserves an epistemic status. First, I will

[2] In this paper, I presuppose monism about validity – that is, that there is a single general sense in which an inference can count as valid or invalid. But this is just for convenience. The central points here will also apply to pluralism about validity.

argue that if that epistemic status is factive, such as *being in a position to know* or *having conclusive evidence for*, then the account immediately runs into some trouble. If, in addition, we endorse certain familiar logical principles – notably, Reductio ad Absurdum and Conditional Proof – then the account runs into still more serious trouble. Second, I will argue that if the relevant epistemic status is non-factive, such as *is rationally committed to* or *has justification for believing*, then a similar problem arises if we accept the familiar logical principles as well as a sufficiently strong "level-bridging" principle. For instance, the proposal concerning commitment runs into trouble given any of the following level-bridging principles: If a thinker is committed to a claim, she is committed to being committed to the claim; If a thinker is committed to being committed to a claim, she is committed to the claim; If a thinker is committed to not being committed to a claim, she is not committed to the claim; A thinker cannot be committed to both a claim and the claim that she is not committed to the claim; A thinker cannot be committed to both the negation of a claim and the claim that she is committed to the claim. Finally, I will argue that closely related problems arise for the most natural characterization of validity in terms of rational credence.

There are ways to avoid these difficulties – for instance, by giving up a familiar logical principle, by rejecting the relevant level-bridging principles, or by moving to an epistemic characterization not in terms of the preservation of some epistemic status. But the difficulties provide a new and, I think, interesting reason to be wary of epistemic characterizations of validity.

In the next section, I'll discuss possible motivations for an epistemic characterization of validity. After that, I'll turn to the difficulties facing such characterizations. Before I do so, let me first make three quick clarifications. First, the difficulties facing epistemic accounts of validity do not rely on the accounts providing *definitions* of validity. It suffices that the accounts provide *characterizations* of validity – that is, conditions that are jointly necessary and sufficient for validity. Second, the difficulties rely on the relevant conditions being both necessary and sufficient. Nothing in what follows poses any difficulty for the project of identifying “bridge principles” that specify the epistemic consequences of validity claims.[3] Finally, the difficulties presented here do not rely on the target accounts making use of an epistemic notion. What they do rely on is that the accounts involve the preservation of some property from the premises to the conclusion, a property that obeys a sufficiently strong level-bridging principle. I focus on epistemic accounts because they are the most natural and most popular non-truth-related views, and because the relevant level-bridging principles are plausible for many epistemic notions. But it is worth keeping in mind that the difficulties here arise for a broader range of views.[4]

2. Motivations for an Epistemic Characterization

What might motivate an epistemic characterization of validity? The most compelling motivation, to my mind, comes from the significance of validity. Validity is important – it matters whether a deductive inference counts as valid or invalid. In what way is validity important? A natural view is that validity is important because it is epistemically significant. Indeed, it is highly plausible that there is a connection between validity and some epistemic status such as knowledge, evidence, justified belief, or rational credence. For instance, perhaps a thinker who is in a position to know the premises of a valid inference is in a position to know the conclusion. Or

[3] The use of the term “bridge principle” is due to MacFarlane (2004), applied to the case of logical validity. See Steinberger (2022) for a recent survey of the literature on bridge principles for logical validity.
[4] For example, the same difficulties will arise for characterizations of validity in terms of the preservation of necessary truth, determinate truth, settled truth (what’s true at a given time in all branches in the open future), and so forth.

perhaps a thinker who has conclusive evidence for the premises has conclusive evidence for the conclusion. Or perhaps a thinker who has justification for believing the conjunction of the premises has justification for believing the conclusion. Or perhaps a thinker who has a high rational credence in the conjunction of the premises has a high rational credence in the conclusion. Or something of the sort.

Principles connecting *logical* validity and epistemic statuses have been widely endorsed. Proposed principles tend to specify the epistemic consequences of logical validity claims. For instance, there are many proponents of "closure principles" for knowledge.[5] These closure principles are aimed at capturing the intuitive thought that "deduction is a way of extending one's knowledge."[6] Notice, however, that the intuitions supporting these principles don't seem to have anything particular to do with logical validity. Analytic or conceptual validity would seem to be exactly on a par.[7]

There are difficulties facing many of the proposed connections between validity and epistemic statuses.[8] But as difficult as it is to formulate a counterexample-free principle, it is also difficult to deny that there is some connection between validity and epistemology. Validity seems relevant to how we ought to reason or what we ought to believe.

This observation can be used to argue for an epistemic characterization of validity. Suppose, for instance, that a thinker who is in a position to know the premises of a valid inference is in a position to know the conclusion. This principle specifies a necessary condition on validity. Given

[5] See, for instance, Williamson (2000, ch. 5) and Hawthorne (2004, ch. 1).
[6] Williamson (2000, p. 117).
[7] Similarly, consider the popular view that rational credences are probabilistically coherent. This view entails a connection between logical validity and rational credence: If an inference is logically valid then the rational credence in the conclusion is at least as high as the rational credence in the conjunction of the premises. Presumably, this connection should be generalized to include analytic or conceptual validity, too.
[8] See Harman (1986) and MacFarlane (2004) for some important difficulties. Harman and MacFarlane focus on logical validity, but their points generalize.

this necessary condition, the corresponding sufficient condition also seems plausible. If the significance of validity consists in the fact that valid inferences preserve what a thinker is in a position to know, it is plausible that whenever an inference preserves what a thinker is in a position to know – or, perhaps, whenever there is a guarantee that an inference preserves what a thinker is in a position to know – the inference counts as valid. This suggests that we endorse an epistemic characterization of validity.

A related line of thought can used to argue for an epistemic *definition* of validity. Suppose again that a thinker who is in a position to know the premises of a valid inference is in a position to know the conclusion. What explains the truth of this principle?[9] It seems unlikely that this connection between validity and knowledge is merely a brute fact. Some kind of explanation of the connection seems called for. A snappy explanation is that validity is defined in terms of the preservation of *being in a position to know*. For an inference to be valid just is, as a matter of definition, for the inference to preserve this epistemic status. And similarly for other potential connections between validity and an epistemic status. What this suggests is that we should endorse an epistemic definition of validity.[10]

---

[9] The challenge of explaining the epistemic significance of validity is a focus of recent work by Prawitz. See, for example, Prawitz (2012). One way of understanding Carroll (1895) – though not, perhaps, what Carroll himself had in mind – is as posing the challenge of explaining the tie between logical validity and epistemology.

[10] Indeed, if validity is defined in proof-theoretic or model-theoretic terms, it is a mystery why validity has epistemic significance. Why should it matter for epistemology what can be deduced in some particular formal proof theory? And why should it matter for epistemology what preserves truth in a model, where models are a kind of algebraic construction? If validity is defined in terms of the preservation of truth, then there may be an answer to the question of why validity is epistemically significant, but there is the danger that this will require an overly externalistic view in epistemology, such as reliabilism about justification with all of its attendant problems.

A related line of thought appears in Field (2015). Field argues that logical validity should not be defined in proof theoretic or model theoretic terms because such definitions do not capture what's at stake in a dispute over which inferences are valid. It can be agreed by all parties that an inference is valid according to some formal proof theory, or is valid in all classical set theoretic models, without there being agreement on whether the inference is genuinely valid. Interestingly, Field's line of thought bears a family resemblance to what Gibbard (2003, pp. 23-9) calls the "what's at issue?" argument for moral non-naturalism in Moore (1903, sec. 11).

Other motivations for an epistemic characterization of validity do not directly support using an epistemic notion in defining validity but instead directly oppose using truth (or a cognate notion) in defining validity.[11] Assuming that validity has a definition (and is not an undefined primitive), it is tempting to think that if the definition of validity does not involve truth, it must involve knowledge, evidence, commitment, justification, or some other epistemic status.

Arguments against defining validity in terms of truth can be divided into two groups – arguments depending on substantive theoretical claims about language or thought and arguments depending on more technical considerations. One argument in the former group stems from inferentialism,[12] the view that the meaning of a word is constituted or determined by[13] inferences involving the word and is not to be explained in terms of representational notions such as truth or satisfaction. A natural way to develop an inferentialist account of meaning is to claim that the inferences that constitute or determine the meaning of a word must all be valid. Since the meaning of a word is not to be explained in terms of representational notions, validity itself should not be explained in terms of representational notions, either.[14] Notice that this motivation doesn't depend on endorsing inferentialism for the entire vocabulary of our language – it suffices that we endorse inferentialism for some sub-vocabulary, such as the logical vocabulary.

A related argument comes from deflationism about truth.[15] Deflationism about truth is the view that (i) truth does not have a substantial nature, (ii) truth cannot play a genuinely explanatory role, and (iii) all that can be said about truth is that the predicate "true" plays a

[11] Defining validity in terms of truth in a model may avoid these difficulties. But given such a view, the problem of accounting for the epistemic significance of validity becomes still more pressing.
[12] See Steinberger and Murzi (2017) for a survey of inferentialism.
[13] There are thus two versions of inferentialism – one that takes inferentialism to be a semantic thesis ("constituted by") and one that takes inferentialism to be a metasemantic thesis ("determined by").
[14] There is also an analogous line of thought that relies on a conceptual-role-based account of mental content rather than an inferentialist account of linguistic content.
[15] See Armour-Garb, Stoljar, and Woodbridge (2021) for a survey of deflationism about truth.

certain expressive role in our language – for instance, in enabling us to indirectly assent to claims and to state generalizations that would otherwise be difficult or impossible to state. Inferentialism and deflationism about truth naturally go together, but the two views do not obviously entail one another. The motivation from deflationism is this: If validity is a substantive notion and truth is not, then presumably truth cannot play the starring role in defining validity. Some other notion must play the central part. Plausibly, that will be an epistemic notion.

There are two arguments against defining validity in terms of truth that are based on more technical considerations. The most natural ways to use truth in defining validity – for instance, defining validity as necessary truth-preservation for some reasonable kind of necessity – entail that valid inferences preserve truth. Both of the technical arguments target this specific claim.[16] Since fully articulating these arguments would require a fair bit of space, I won't describe them in any real detail here, but merely provide high-level summaries.[17]

The first technical argument concerns the Curry paradox. This is the paradox concerning a sentence S that is equivalent to the claim that if S is true then ⊥, where ⊥ is some absurdity. There are several ways to use this sentence to argue to the absurd conclusion ⊥. The claim that valid inferences preserve truth (i) blocks common diagnoses of the most familiar version of the Curry paradox (such as the diagnoses that the problem is due to the use of Conditional Proof or that the problem is due to the use of the inference rule that permits inferring from any claim A to

[16] Field (2008, ch. 19.2) tries to explain away the intuitive appeal of the claim that valid inferences preserve truth. He notes that there is a natural argument for this claim relying on Modus Ponens, Conditional Proof, and intuitive rules for the truth predicate. But, Field points out, these resources also suffice for generating a Curry paradox (assuming the usual structural rules). So something the natural argument relies upon must be rejected. Field (2015, p. 39) also suggests that for arguments involving only "ordinary" sentences (e.g., not involving self-reference or infinite chains of reference), validity coincides with necessary truth preservation.

[17] See Murzi and Shapiro (2015) for a discussion of these arguments and a defense of the claim that valid arguments preserve truth.

the claim that A is true),[18] and also (ii) gives rise to a particularly simple and hard-to-avoid version of the Curry paradox.[19]

The second technical argument concerns Gödel's second incompleteness theorem. It seems possible to construct a reasonably powerful mathematical or semantic theory that has axioms we accept as true and rules of inference we accept as valid. Given this, it seems that we should be able to consistently extend such a theory to include the claims that its axioms are true and its rules are valid. If valid inferences preserve truth, this principle presumably can be consistently added to the theory, too. But then, there will be a simple inductive proof that can be carried out in the theory itself to the conclusion that all theorems of the theory are true. Since truth entails consistency, this procedure seems to yield a recursively axiomatizable theory that is reasonably powerful (e.g., is stronger than Peano Arithmetic), is consistent, and can prove itself consistent, violating Gödel's second incompleteness theorem. A natural diagnosis of the problem is that it stems from the claim that validity preserves truth.[20]

These motivations are not fully conclusive. In response to the first motivation, concerning the importance of validity, even assuming that validity is epistemically significant, one could reject the claim that the relevant epistemic condition provides a sufficient condition for validity. One could also provide a different explanation of the connection between validity and an epistemic status. In response to the motivations concerning inferentialism and deflationism about truth, one could reject those views or try to argue that they are in fact compatible with defining (or characterizing) validity in terms of truth-preservation. In response to the two technical arguments, one could find a different diagnosis of the Curry Paradox or a different way to block

[18] See Field (2015, pp. 39-41).
[19] See Beall (2009, p. 35). See Field (2008, pp. 377-8) for a related argument in a slightly different context.
[20] This is a variant of an argument due to Field (2006; 2008, chs. 12.4 and 19.3).

the apparent violation of Gödel's second incompleteness theorem. And even if we were to conclude from the theoretical or technical motivations that validity should not be defined in terms of truth, this doesn't preclude us from providing a different kind of non-epistemic definition of validity or from treating validity as an undefined primitive.[21] Nevertheless, taken together, these motivations provide a surprisingly strong case against defining validity in terms of truth and in favor of defining validity in epistemic terms.

3. Factive Operators

The previous section provided motivations for accepting an epistemic characterization – and indeed, an epistemic definition – of validity. Let me now turn to the difficulties facing epistemic characterizations. In this section, I'll consider the proposal that an inference is valid just in case the inference preserves some factive epistemic status.

3.1. Knowledge

Let's start with the idea that validity should be characterized in terms of knowledge. This is a natural suggestion for at least two reasons. First, as mentioned above, closure principles for knowledge are widely accepted. The idea that validity should be characterized in terms of knowledge is essentially the idea that such a principle should be strengthened from a conditional to a biconditional. Second and relatedly, a view that is popular with some epistemologists is "knowledge-first epistemology", which is (roughly) the view that (i) knowledge is the primary epistemic notion, (ii) other doxastic and epistemic notions (such as justification) should be

[21] Interestingly, analogues of the motivation from inferentialism and the two technical considerations arise against some epistemic definitions of validity, too. This is clearest for definitions in terms of the preservation of a factive epistemic status such as knowledge. But there are analogous motivations against definitions in terms of the preservation of some non-factive statuses, too.

defined in terms of knowledge, and (iii) the epistemic norms governing our practices (such as the norms governing assertion, reasoning, and belief) involve knowledge.[22] If one is a knowledge-first epistemologist and thinks that validity should be defined in epistemic terms, it will be tempting to define validity in terms of knowledge.

Somewhat surprisingly, I don't know of many proposals of this sort. The only one I'm aware of in the literature is due to Sundholm:

> What does it mean for [a general inference form] I to be valid? We consider how an inference according to I is used. In such use one takes it for granted that the premisses J_1, ..., J_k are known and goes on to obtain knowledge of J. Thus, under the epistemic assumption that the judgements J_1, ..., J_k are all known, one has to make the judgement J known. Sundholm (2012, p. 950)[23]

But I think that, at least for those philosophers that treat knowledge as the central epistemic notion, characterizing validity in terms of knowledge should be attractive.

It will prove useful to be a bit more precise about just what an account of validity in terms of the preservation of knowledge should look like. A first stab of such an account is that the

[22] See Williamson (2000). Also see McGlynn (2014) for critical discussion.

[23] Sundholm's account faces difficulties that are independent of the central issues raised in this paper. Since Sundholm requires that in a valid inference, one "obtains" knowledge of the conclusion from knowledge of the premises, Sundholm's account has the consequence that validity is non-reflexive. Sundholm might respond by claiming that there is no inference – movement in thought or language – from a judgment to itself, so that is an acceptable consequence of his account. But the problem goes deeper since it is plausible that there are claims A and B such that the inference from A to B and the inference from B to A are both valid. (The DeMorgan laws provide plausible examples of this.) Sundholm's account would seem to rule this out. More generally, the problem is that there are valid inferences from premises one knows to a conclusion one knows, where one's knowledge of the conclusion does not depend on the knowledge of the premise – indeed, the flow of knowledge may be the other way around. An improved version of Sundholm's account would require only that in a valid inference if one knows the premises and carries out the inference one knows the conclusion, and omit anything to do with obtaining knowledge of the conclusion from knowledge of the premises. Thanks to an anonymous referee for raising this issue.

inference from A to C is valid just in case for every thinker, if the thinker knows that A then she knows that C.

This proposal is too crude for at least two reasons. There is a problem concerning sufficiency and a problem concerning necessity. The problem concerning sufficiency is that it could have been the case that every actual thinker knows both that A and that C, for some claims A and C where the inference from A to C is not valid. In such a scenario, if a thinker knows that A then she knows that C.[24] So the inference from A to C will fit the proposed characterization but fail to be valid. The problem concerning necessity is that a thinker can know some claim A and fail to know a simple consequence of A – such as the disjunction either A or B – simply because the thinker does not believe the consequence (or because the thinker believes it but on bad grounds). So an inference may be valid but fail to fit the proposed characterization.

In response to the problem with sufficiency, the natural suggestion is to strengthen (or otherwise modify) the quantified conditional "for every thinker, if the thinker knows that A then she knows that C". There are several ways in which this could be done. For instance, the quantified conditional could be required to be necessary on some reasonable sense of "necessary", such as conceptual necessity. Or it could be required to be a priori. Or it could be generalized to require that if a thinker knows that A′ then she knows that C′ for all appropriate variants <A′,C′> of <A,C>. Or it could be required that in every model (satisfying appropriate constraints) in which a thinker knows that A, the thinker knows that C. Or something similar.[25] In response to the problem with necessity, the natural suggestion is to characterize validity not in

[24] If there is a constraint on conditionals that the antecedent be relevant to the consequent, A and C should be chosen to satisfy the appropriate relevance constraint.

[25] An alternative suggestion is to appeal to an explanatory relation, such as metaphysical ground. The idea would be that the fact that a thinker knows that C is explained by the fact that the thinker knows that A. This suggestion is appealing but it faces a version of the problem that was raised in footnote 23. Explanatory relations are typically taken to be irreflexive and, more generally, asymmetric. But it is plausible that validity is reflexive. Even if it is not, it is plausible that validity is not asymmetric.

terms of what a thinker knows but in terms of what the thinker is in a position to know (given her epistemic position), or what is in principle knowable, or something of the sort. The precise details of how best to strengthen (or modify) the conditional and which knowledge-like operator should replace knowledge will not matter in what follows.

We can symbolize the resulting account of validity as follows:

$$A \vdash C \text{ just in case } KA \vdash_{base} KC$$

Here, "$A \vdash C$" is the claim that the inference from A to C is valid. "K" is the relevant knowledge-like factive operator. For instance, it could be the operator "it is knowable that". Finally, "$\vdash_{base}$" is an operator that replaces the indicative conditional. For instance, "$A \vdash_{base} C$" could be the claim that it is conceptually necessary that if A then C.[26, 27]

It will turn out to be useful to allow multiple claims on the left-hand side of "$\vdash_{base}$". One way to do so would be to understand the claims on the left-hand side of "$\vdash_{base}$" to be implicitly conjoined. On this suggestion, for instance, "$A, B \vdash_{base} C$" could be the claim that it is conceptually necessary that if both A and B then C. But we could instead take "$\vdash_{base}$" to be an operator that directly applies to a plurality of claims on the left and a single claim on the right.

There are two natural ways of extending the account of validity from single-premise to multi-premise inferences. The first is that K is preserved from each of the premises to the conclusion.

[26] We can even take "$A \vdash_{base} C$" to be $\vdash A \rightarrow C$, or even $A \vdash C$. In the first case, we'd have to either (i) provide an antecedent account of when a claim is a validity and treat the characterization as showing how to extend that account to valid inferences or (ii) treat the proposed account as providing a constraint on validity rather than a definition of validity. In the second case, we'd have to either (i) provide an antecedent account of when there is a valid inference from KA to KC and treat the characterization as showing how to extend that account to all valid inferences or (ii) treat the proposed account as providing a constraint on validity rather than a definition of validity.

[27] In the literature, the use of an operator like $\vdash_{base}$ shows up in proposed epistemic characterizations of validity. (Typically, the operator is not further explained.) See, for example, the quote by Prawitz in section 3.2 below, which uses the phrase "follows that" and the quote by Brandom in section 4.1. below, which uses the phrase "as a consequence of".

The second is that K is preserved from the conjunction of the premises to the conclusion. In symbols:

$$A_1, \ldots, A_n \vdash C \text{ just in case } KA_1, \ldots, KA_n \vdash_{\text{base}} KC$$

or

$$A_1, \ldots, A_n \vdash C \text{ just in case } K(A_1 \,\&\, \ldots \,\&\, A_n) \vdash_{\text{base}} KC$$

In what follows, all that will matter is the single-premise case. So the difference between these two extensions will not be relevant.[28]

Now that we have the proposal in a bit more focus, we can turn to the difficulty facing it. Consider the claim A&¬KA. This is analogous to a Moore-paradoxical sentence, replacing belief with the operator K.[29] Given the characterization of validity in terms of K, we can show that $A\&\neg KA \vdash \bot$, where $\bot$ is an absurdity. The argument for this is straightforward. We show that $K(A\&\neg KA) \vdash_{\text{base}} K\bot$, using reasoning reminiscent of that of Fitch's paradox of knowability.[30] This part of the argument relies on the left-to-right direction of the proposed characterization of validity. It then follows from the right-to-left direction of the proposed characterization of validity that $A\&\neg KA \vdash \bot$.

In more detail, the proof relies on the following principles:

---

[28] An advantage of the former extension is that it better captures the idea that valid inference preserves knowledge from the premises (individually) to the conclusion. A disadvantage of the former extension is that it may not yield $A, B \vdash A\&B$ due to a variant of Makinson's (1965) preface paradox. A disadvantage of the latter extension is that it does not cohere with an inferentialist treatment of conjunction, at least when taken to be a definition. The reason is that on an inferentialist account of "&", the meaning of "&" is constituted or determined by valid inferences, presumably including multiple-premise inferences. But the proposed definition of validity for multiple-premise inferences presupposes that "&" is already assigned a meaning, so there is a problem of circularity. See Peacocke (1992, p. 180) for discussion of this kind of issue.

[29] Moore's paradox first appeared in Moore (1942, pp. 542-3).

[30] See Fitch (1963). It is now known that the argument is originally due to Alonzo Church, appearing in a referee report of an earlier paper by Fitch.

(i) Conjunction Elimination for $\vdash$: $A\&B \vdash A$ and $A\&B \vdash B$

(ii) Factivity of K for $\vdash_{base}$: $KA \vdash_{base} A$

(iii) Explosion for $\vdash_{base}$: $A, \neg A \vdash_{base} B$

(iv) Cut for $\vdash_{base}$: If both $A, D \vdash_{base} C$ and $B \vdash_{base} D$ then $A, B \vdash_{base} C$; and if both $D, B \vdash_{base} C$ and $A \vdash_{base} D$ then $A, B \vdash_{base} C$

(v) Contraction for $\vdash_{base}$: If $A, A \vdash_{base} C$ then $A \vdash_{base} C$

These principles are all natural for the construals of $\vdash_{base}$ considered above.

The proof, then, goes as follows: By Conjunction Elimination, $A\&\neg KA \vdash A$ and $A\&\neg KA \vdash \neg KA$. By the characterization of validity, $K(A\&\neg KA) \vdash_{base} KA$ and $K(A\&\neg KA) \vdash_{base} K\neg KA$. By Factivity, $K\neg KA \vdash_{base} \neg KA$. By Explosion, $KA, \neg KA \vdash_{base} K\bot$. By three uses of Cut, $K(A\&\neg KA), K(A\&\neg KA) \vdash_{base} K\bot$. By Contraction, $K(A\&\neg KA) \vdash_{base} K\bot$. By the characterization of validity, $A\&\neg KA \vdash \bot$.[31]

This is an awkward result. Suppose "K" stands for "it is knowable that". Presumably, it is genuinely possible that there is a truth that is not knowable. For instance, plausibly it is genuinely possible that it is true but unknowable that there is an even number of stars in the universe, or that the happiest person who ever lived was alive in the year 5,437 BCE. And if it is genuinely possible that there is a true but unknowable claim, then presumably the conjunction of A with the claim that A is not knowable is not always inconsistent. And similarly for other candidates for "K".

One might respond to this result by distinguishing between entailment, understood as an abstract relation between claims, and validity, understood as applying to potential inferences in

[31] If we accept the first of the two extensions of the proposed characterization of validity to the multi-premise case, we can show $A, \neg KA \vdash \bot$ without relying on Conjunction Elimination or Contraction: By Factivity, $K\neg KA \vdash_{base} \neg KA$. By Explosion, $KA, \neg KA \vdash_{base} K\bot$. By Cut, $KA, K\neg KA \vdash_{base} K\bot$. So $A, \neg KA \vdash \bot$.

language or thought. One might claim that while the conjunction of A and ¬KA does not entail ⊥, the inference from the conjunction to ⊥ is valid. The idea is that, just as the claim "A but I don't know that A" cannot rationally be accepted, neither can "A&¬KA". One way to represent this is to say that the inference from A&¬KA to ⊥ is valid.

This response has some surface plausibility, but it faces serious difficulties. One problem is that, even if one can't rationally outright accept A&¬KA, one can still suppose that A&¬KA and reason about what is the case under that supposition. For many such suppositions, ⊥ is not the case. For instance, supposing that there is an even number of stars in the universe and that this claim is unknowable, we can validly infer many consequences – for instance, that there is not an odd number of stars – but we cannot validly infer an absurdity. It is plausible that the account of valid inference under suppositions should exactly parallel the account of inference from claims that are outright accepted. So there is reason to avoid saying that one can validly infer from A&¬KA to ⊥.

There is a second problem. Given some natural principles governing ⊢, the result that A&¬KA ⊢ ⊥ has further, unpleasant consequences. By Reductio ad Absurdum, A ⊢ ¬¬ KA. (Here, I'm using Reductio ad Absurdum in the following form: if A&B ⊢ ⊥ then A ⊢ ¬B.) Then by Double Negation Elimination and Transitivity, A ⊢ KA. Finally, by Conditional Proof, ⊢ A→KA. In other words, it is a validity that if A then KA.[32] If, for instance, "KA" says that A is knowable, then it is a validity that if A then A is knowable. Since presumably it is not true that for every claim A, if A then A is knowable, it is highly problematic that this turns out to be a validity.[33]

[32] In a logic without Double Negation Elimination, such as Intuitionistic logic, there is still a problem. We can no longer show that ⊢ A→KA, but we can show that ⊢ A→¬¬KA, which is nearly as bad.
[33] An anonymous referee asks whether this result still goes through if "$\vdash_{base}$" stands for a counterfactual conditional. On the standard Lewis-Stalnaker variably-strict semantics for counterfactual conditionals, the proof in the text that

We can go further. By the factivity of K, $\vdash KA \rightarrow A$. So we can show that $\vdash A \leftrightarrow KA$. That is, the distinction between "A" and "KA" seems to disappear. Moreover, if we have some device of generalization, such as quantifying into sentence position, we can go still further and show that $\vdash \forall \mathbb{A}(\mathbb{A} \leftrightarrow K\mathbb{A})$. That is, it is a validity that for any sentence, it is true just in case it is K. (If we don't have a direct way to quantify into sentence position, we can get the same effect using a truth predicate.) But even if we don't have a device of this sort, it is extremely awkward to accept for each A, $\vdash A \leftrightarrow KA$. It is extremely implausible that $A \leftrightarrow KA$ is a validity for every A.

How might one respond to this result while holding on to the account of validity? One option is to try to endorse the conclusion and claim that $A \leftrightarrow KA$ is in fact a validity. For most choices of "K", such as what a specific thinker is in a position to know given her epistemic position, this is grossly implausible. If "K" is understood as what is in principle knowable, then this conclusion may be (slightly) better. But it still would require endorsing a strong kind of antirealism.[34]

The only other apparent option is to reject one or more of the principles at issue in the argument, most plausibly Conditional Proof or Reductio ad Absurdum.[35] The trouble with this suggestion is that these are highly plausible metarules. Moreover, anyone who endorses an epistemic characterization of validity because they endorse inferentialism for the logical connectives will run into difficulty in rejecting one of these metarules. The reason is that

---

$A\&\neg KA \vdash \bot$ no longer goes through, since Cut fails. However, it is straightforward to show $K(A\&\neg KA) \vdash_{base} K\bot$ directly (assuming K distributes over conjunction) and so $A\&\neg KA \vdash \bot$ still obtains. What fails are the Reductio ad Absurdum and Conditional Proof steps. In any event, this is not a promising approach to take for several reasons: First, $A\&\neg KA \vdash \bot$ is bad enough. Second, Reductio ad Absurdum and Conditional Proof are intuitive, and it is awkward to be forced to reject them. Finally, counterfactual accounts typically fail due to versions of the conditional fallacy.

[34] Moreover, this is where one can appeal to Fitch's (1963) paradox to argue that the claim that all truths are knowable entails the far more implausible claim that all truths are known.

[35] There are other possible culprits. Advocates of a paraconsistent logic will reject Explosion. Advocates of a substructural logic will reject Cut or Contraction. I view these diagnoses as more radical than the rejection of Conditional Proof or Reductio ad Absurdum. Moreover, the usual reasons to reject these principles (e.g., the semantic paradoxes) don't seem to be relevant to the issue at hand. So proponents of such views may end up saddled with $\vdash A \leftrightarrow KA$ restricted to "ordinary" claims. (The rejection of Contraction also doesn't enable one to fully avoid the problem, since there will still be a way to show $\vdash A\&A \rightarrow KA$.)

inferentialists about logic tend to view Conditional Proof as constitutive of the meaning of the conditional and Reductio ad Absurdum as constitutive of the meaning of negation. So they will be loath to reject these metarules.

Even those who reject Conditional Proof (for instance, on grounds relating to the semantic paradoxes or the paradoxes of vagueness) typically endorse some weakening of it, perhaps to claims that do not involve self-reference or non-well-founded chains of reference, or to claims that don't involve a determinacy operator. Presumably, such a weakened version of Conditional Proof will still apply to everyday claims about the world as well as to claims to the effect that an everyday claim about the world is knowable. So even those who reject Conditional Proof will be forced to accept $\vdash A \rightarrow KA$ for "ordinary" claims, which is just as problematic. (And similarly for Reductio ad Absurdum.)[36]

What all of this suggests is that a characterization of valid inference in terms of the preservation of *being in a position to know* or *is knowable* or something similar simply does not work.

3.2. Conclusive Evidence

Some philosophers have characterized validity, or a related notion, by appealing to a factive epistemic state other than knowledge. Here, for example, is Prawitz:

[36] A different suggestion is that the problem is due to the fact that $A\&\neg KA$ uses the same operator that appears in the characterization of validity. So perhaps we should modify the definition to avoid this. The natural suggestion to make is to move to a hierarchy of K-operators, $K_1, \ldots, K_n, \ldots$, on the model of Tarski's theory of truth. The idea would be to characterize validity something like: $A \vdash C$ just in case $K_nA \vdash_{base} K_nC$, where n is one greater than the highest K-subscript appearing in A and C (or in any claim directly or indirectly referred to by A or C). The idea of typing knowledge has been proposed in other contexts, including as a response to the knower paradox, such as in Anderson (1983) and, closer to the issues discussed here, as a response to Fitch's paradox, such as in Paseau (2008). In the present context, however, this suggestion won't help. The reason is that if A is K-free, K_2A will presumably be equivalent to K_1A, and so the problematic reasoning will still go through. Thanks to Bruno Whittle for suggesting both this approach and the problem with it.

> In sum, it is required of a proof that all its inferences are successful. To characterize a proof as a chain of inferences, as we usually do, we thus need this notion of successful inference. It is convenient to have a term for this, that is, for inferences that can be used legitimately in a proof, and I have called them *legitimate inferences*…. Accordingly, a generic inference is said to be legitimate, if a subject who makes the inference and has evidence for its premisses thereby gets evidence for the conclusion; or more precisely, it should follow that she has evidence for the conclusion from the assumptions that she performs the inference and has evidence for the premisses. We can now say that a deductive proof is a chain of legitimate inferences. Prawitz (2015, p. 73)[37]

By "evidence" Prawitz means conclusive evidence, which is most naturally understood to be factive.[38] Let E be a one-place sentential operator such that "EA" stands for the claim that the relevant thinker possesses conclusive evidence for A. Prawitz's discussion suggests that we characterize legitimate inference, $\vdash$, as follows, at least for the single premise case:[39]

$A \vdash C$ just in case $EA \vdash_{base} EC$

---

[37] Notice that Prawitz's initial gloss of "legitimate inference" faces the same difficulty that was raised in footnote 23 due to his use of the word "thereby". Prawitz's subsequent gloss (after the "more precisely") avoids this difficulty.

Prawitz does not identify legitimacy with validity but presents a more complicated definition of "valid inference". In particular, Prawitz defines an inference to be valid if the operation involved in the inference is such that when applied to conclusive grounds for the premises it yields a conclusive ground for the conclusion. This definition also seems to face a version of the difficulty developed below, but the issues are a bit more delicate here due to the appeal to an operation.

[38] If "conclusive evidence" is interpreted to mean indefeasible evidence, there is a worry that there isn't much if any conclusive evidence. The force of (almost) any evidence can be defeated by higher-order considerations. For instance, one might think that the claim A provides conclusive evidence for the claim either A or B. But someone could acquire strong but misleading evidence that they were suffering from a condition in which they mistook conjunctions for disjunctions, or that they were unreliable in inferences concerning the particular claim A. This could defeat the evidential support that A provides to either A or B. To rule out such cases, we presumably should define "conclusive evidence" so that the prospect of higher-order evidence isn't relevant to whether evidence counts as conclusive.

It may be preferable to replace "conclusive evidence" with "conclusive grounds". This is because for some claims, talk of evidence seems out of place. Examples include simple logical, mathematical, and conceptual truths. It is plausible that beliefs in such claims are typically not justified by evidence but in some other way.

[39] Prawitz presumably would also require the right-hand side to include the claim that the relevant subject drew the inference. This slightly complicates matters but doesn't avoid the problem raised here.

The trouble with this proposal is the same as for the proposal involving knowledge. Since having conclusive evidence is factive, we can show $A\&\neg EA \vdash \bot$. Moreover, by the same reasoning as above, we can show $\vdash A\rightarrow EA$, and then $\vdash A\leftrightarrow EA$, and perhaps even $\vdash \forall\mathbb{A}(\mathbb{A}\leftrightarrow E\mathbb{A})$. That is, it is a validity that for every claim, it is true just in case the relevant subject has conclusive evidence for it. This is highly implausible. One can avoid the result that $\vdash A\rightarrow EA$ by giving up Conditional Proof, Reductio ad Absurdum, or some other natural principle. But that is a real cost, and even if one does so, one is still left with $A\&\neg EA \vdash \bot$.

3.3. Factive Operators in General

As should be apparent, the argument generalizes. Suppose O is any one-place factive operator. Suppose we characterize $\vdash$ as follows:

$A \vdash C$ just in case $OA \vdash_{base} OC$

Then, if we have Conjunction Elimination for $\vdash$ as well as Explosion, Cut, and Contraction for $\vdash_{base}$, we can readily show that $A\&\neg OA \vdash \bot$. We can further show that if, in addition, $\vdash$ obeys Conditional Proof, Reductio ad Absurdum, Double Negation Elimination, and Transitivity, $\vdash A\rightarrow OA$. We can then show $\vdash A\leftrightarrow OA$, and perhaps even $\vdash \forall\mathbb{A}(\mathbb{A}\leftrightarrow O\mathbb{A})$. These results obtain whether O stands for an epistemic operator such as "is knowable" or "is provable" or a non-epistemic operator such as "is determinately true" or "is necessarily true". What these results suggest is that we had better not try to characterize valid inference as the preservation of a property that (i) is factive and (ii) is "strong" in the sense that $A\rightarrow OA$ is not always a validity.[40]

[40] The problem may even arise for characterizations of validity in terms of the preservation of truth, since on many formal theories of truth $A\rightarrow Tr(\langle A\rangle)$ is not always a validity for reasons to do with the semantic paradoxes. Of course, proponents of such theories of truth tend to reject Conditional Proof.

4. Non-Factive Operators

4.1. Commitment

What about characterizing validity in terms of the preservation of a non-factive operator? Brandom, for instance, characterizes an important form of (correct) inference – which he takes to be a generalization of (correct) deductive inference – in terms of the preservation of a kind of commitment:

> Inheritance of commitment (being committed to one claim as a consequence of commitment to another) is what will be called a *committive*, or commitment-preserving inferential relation. Deductive, logically good inferences exploit relations of this genus. But so do materially good inferences, such as A is to the West of B, so B is to the East of A.… Brandom (1994, p. 168, italics in the original)

Similarly, Incurvati and Schlöder suggest that an important form of inference can be characterized in terms of the preservation of commitments:

> Our inference rules are about the necessary consequences of commitments with respect to coherence. That is, given that a speaker has undertaken certain commitments (displayed some attitudes), the rules tell us what further commitments that speaker is bound by on pain of incoherence. Since we linked commitment to coherence, this means that our inference rules preserve commitment. Incurvati and Schlöder (2017, p. 748)[41]

[41] This is not meant to be merely a sufficient condition but a definition. Incurvati and Schlöder (forthcoming, p. 12) quote the second sentence of this passage and explicitly state that this is how they "define commitment-preserving inference".

Suppose "MA" stands for the claim that the relevant subject is committed to A.[42] Suppose we characterize validity (or a closely related notion), $\vdash$, as follows:

$A \vdash C$ just in case $MA \vdash_{base} MC$

Does this suggestion run into difficulty, too?[43]

Commitment is not factive – being committed to a claim doesn't entail that the claim is true. So we cannot simply apply the results concerning factive operators to this case. However, there are similar worries that arise here, as well.

Consider the following two principles. First, there is the modal 4 axiom: (4) $\vdash_{base}$ $MA \rightarrow MMA$. If one is committed to a claim, then one is committed to being committed to the claim. Second, there is the converse of the modal 5 axiom: (5_c) $\vdash_{base} M\neg MA \rightarrow \neg MA$. If one is committed to not being committed to a claim, then one is not committed to the claim. This is a weakening of factivity.

Both principles are "level-bridging" principles. They relate higher-order and lower-order commitments. The 4 principle is plausible for commitment. There is something incoherent-seeming about being committed to a claim and not being committed to that commitment. This is especially true if, like Brandom or Incurvati and Schlöder, we understand commitment in terms of what is publicly undertaken via the assertion of a sentence.

Depending on one's conception of commitment, the 5_c principle may be plausible, too. On one conception, one's commitments are purely additive. That is, if one has some commitments,

[42] Is commitment an epistemic notion? The answer seems to be yes for Incurvati and Schlöder's notion since they emphasize the tie to (rational) incoherence. The answer seems plausibly yes for Brandom's notion, too, since commitment to the claim that A is to the West of B *rationally* commits one to the claim that B is to the East of A. More generally, Brandom's notion of committive inference is supposed to play the role of something like analytic or a priori inference, so presumably is intended to be characterized in epistemic terms.

[43] Another potential target is Stalnaker (1975), who defines an inference to be (pragmatically) "reasonable" if it is impossible to accept the premises without being committed to conclusion.

one may thereby incur further commitments (e.g., if one is committed to a conjunction one is thereby committed to its conjuncts). But having some commitments cannot rule out having other commitments. On this conception of commitment, 5_c is not a plausible principle.

On a different conception of commitment, having some commitments can rule out having other commitments, when there is some kind of rational clash between them (e.g., if one is rationally committed to a claim then one cannot rationally be committed to its negation, at least putting aside paraconsistency). On this conception of commitment, 5c is a plausible principle – there is something incoherent-seeming about being committed to not having a commitment while simultaneously having the commitment.

If commitment obeys either of the 4 or 5_c principle, the proposed account of $\vdash$ faces difficulties. Consider first the 4 axiom, $\vdash_{base}$ MA→MMA. By Modus Ponens and Cut, MA $\vdash_{base}$ MMA.[44] So A $\vdash$ MA. This is already implausible – it does not in general follow from a claim that one is committed to the claim. (One way to see this is to note that the account of valid inference under suppositions should exactly parallel the account of inference from outright accepted claims. But it is not plausible that under the supposition of A, one is always committed to A.) Moreover, if $\vdash$ obeys Conditional Proof, $\vdash$ A→MA. That is, it is a validity that if A then one is committed to A. If, in addition, we can quantify into sentence position, we can show that $\vdash \forall\mathbb{A}(\mathbb{A}\rightarrow M\mathbb{A})$. That is, it is a validity that one is committed to every truth. That's a highly unpalatable claim to endorse.

[44] The version of Cut needed here is Cut for theorems: If both A,B $\vdash_{base}$ C and $\vdash_{base}$ B then A $\vdash_{base}$ C. In more recent work, Brandom (2018) suggests that we reject Cut. I'm not sure whether he would reject this weak version of Cut. Even if he would, it would be a further step to reject MA $\vdash_{base}$ MMA, which is simply the rule form of the 4 principle for commitment.

Next, consider the 5_c axiom. We can show that $A \& \neg MA \vdash \bot$ in just the same way we did for factive operators. (Factivity's only role in the argument was to yield 5_c.)[45] And, following the same line of thought as before, if $\vdash$ also obeys Conditional Proof, Reductio ad Absurdum, Double Negation Elimination, and Transitivity, $\vdash A \rightarrow MA$. If, in addition, we have a device for quantifying into sentence position, we can show $\vdash \forall \mathbb{A}(\mathbb{A} \rightarrow M\mathbb{A})$. Again, this is a highly unpalatable claim to endorse.

There is an additional problem. Consider the converse of the modal 4 axiom: (4_c) $\vdash_{base} MMA \rightarrow MA$. If one is committed to being committed to a claim, then one is committed to the claim. This is another plausible principle governing commitment, on either conception of commitment. By Modus Ponens and Cut, $MMA \vdash_{base} MA$. So $MA \vdash A$. Moreover, if $\vdash$ obeys Conditional Proof, $\vdash MA \rightarrow A$. That is, it is a validity that if one is committed to A then A. This is problematic since commitment is not a factive notion. If, in addition, we can quantify into sentence position, we can show that $\vdash \forall \mathbb{A}(M\mathbb{A} \rightarrow \mathbb{A})$. That is, it is a validity that all of one's commitments are true. That's another highly unpalatable claim to endorse.

4.2. Non-Factive Operators in General

Again, these arguments generalize. For instance, the proposal that validity should be characterized in terms of the preservation of having justification to believe will run into the very same trouble. So long as justification obeys the 4 or 5_c axioms – that is, if one has justification to believe a claim one has justification to believe that one has justification to believe the claim, and if one has justification to believe that one does not have justification to believe a claim then one does not have justification to believe the claim – then, given a few natural logical principles, it

[45] This point is originally due to Mackie (1980), in his discussion of Fitch's paradox.

will be a validity that one has justification to believe any truth. And if justification obeys the 4_c axiom – that is, if one has justification to believe that one has justification to believe a claim then one has justification to believe the claim – then, given a few natural logical principles, it will be a validity that one only has justification to believe true claims. Since thinkers do not have justification for believing every true claim, and since thinkers do not have justification for believing only true claims, these are problematic results. And the same will be true for other proposals, such as the proposal that validity should be characterized in terms of the preservation of warranted assertibility.[46]

Very generally, what we've shown is that we had better not try to characterize validity as the preservation of an operator O that either (i) obeys the 4 or 5_c axioms and is strong in the sense that $A \rightarrow OA$ is not always a validity, or (ii) obeys the 4_c axiom and is weak in the sense $OA \rightarrow A$ is not always a validity. Or, at least, this is true given a few natural logical principles. (In the next section, we'll strengthen these results a bit further.)[47]

5. Level-Bridging Principles

The difficulties for epistemic characterizations of validity crucially depend on two ingredients: (i) various logical principles, and (ii) the epistemic notion at issue obeying a sufficiently strong level-bridging principle, such as 4, 5_c, or 4_c. (In the case of a factive operator such as knowledge,

[46] Such a characterization might be inspired by Dummett (1991, pp. 175-6 and 193), who argues that (correct) deductive inference preserves warranted assertibility, and Tennant (1997, p. 262) who writes, "For the anti-realist [such as Tennant himself], the role of logical inference is to preserve warranted assertibility." See Tennant (1987, ch. 13) for further details of his view.

[47] It is also worth mentioning that even if $A \rightarrow OA$ is not immediately implausible for some operator O (e.g., "is in principle justifiable"), a version of Fitch's (1963) paradox may arise that leads to a much more implausible conclusion ("is in fact justified"). See San (2020) for discussion of how versions of Fitch's paradox can arise very generally, without the need for a factive operator.

the level-bridging principle is factivity itself.) This raises a question: How weak can these logical and level-bridging principles be while still generating a problem?

Let's start by considering the logical principles. How weak can they be while still generating the difficulty? Recall that the results in the previous section concerning the 4 and 4_c axioms rely upon Modus Ponens, Cut, and Conditional Proof. Notice, however, that Modus Ponens and Cut are only needed to move from 4 to its rule version $OA \vdash_{base} OOA$, and to move from 4_c to its rule version $OOA \vdash_{base} OA$. So if we have either of these rules to begin with, all we need is Conditional Proof. In other words, we have the following simple result:

Proposition 1. Suppose $A \vdash C$ just in case $OA \vdash_{base} OC$. Suppose $\vdash$ obeys Conditional Proof. Then (i) if $OA \vdash_{base} OOA$ then $\vdash A \rightarrow OA$, and (ii) if $OOA \vdash_{base} OA$ then $\vdash OA \rightarrow A$.

Proof: (i) Suppose $OA \vdash_{base} OOA$. By the characterization of $\vdash$, $A \vdash OA$. By Conditional Proof, $\vdash A \rightarrow OA$. (ii) Suppose $OOA \vdash_{base} OA$. By the characterization of $\vdash$, $OA \vdash A$. By Conditional Proof, $\vdash OA \rightarrow A$.

What about the level-bridging principles? How weak can they be while still generating the difficulty? Consider the principle $\vdash_{base} \neg O(A \& \neg OA)$. If, for instance, "O" stands for the relevant thinker having justification to believe, this principle says that it is not the case that the thinker has justification to believe the conjunction of a claim A with the claim that the thinker does not have justification to believe A. In other words, there is a kind of epistemic clash between believing something and believing that one does not have justification to believe it, a clash that precludes a thinker from having justification to believe the conjunction. Similarly, consider the principle $\vdash_{base} \neg O(\neg A \& OA)$. If "O" again stands for the relevant thinker having justification to believe, this principle says that it is not the case that the thinker has justification to believe the conjunction of the negation of a claim A with the claim that the thinker has justification to believe A. In other words, there is a kind of epistemic clash between believing that one has justification to believe a claim and believing the negation of the claim, a clash that precludes a thinker from having justification to believe the conjunction.[48]

These two principles are intuitively plausible for justification (and for commitment, and for many other epistemic notions). Using reasoning resembling the reasoning above, the former principle can be used to show $\vdash A \rightarrow OA$ and the latter principle can be used to show $\vdash OA \rightarrow A$. (These results depend on Explosion and Cut for $\vdash_{base}$ as well as Reductio ad Absurdum, Conditional Proof, Double Negation Elimination, and Transitivity for $\vdash$.)

We can weaken these level-bridging principles still further. Consider the principles $O(A \& \neg OA) \vdash_{base} O\bot$ and $O(\neg A \& OA) \vdash_{base} O\bot$. It is straightforward to show that, given a few natural logical principles, they yield $\vdash A \rightarrow OA$ and $\vdash OA \rightarrow A$, respectively:

> **Proposition 2.** Suppose $A \vdash C$ just in case $OA \vdash_{base} OC$. Suppose $\vdash$ obeys Reductio ad Absurdum, Conditional Proof, Double Negation Elimination, and Transitivity. Then (i) if $O(A \& \neg OA) \vdash_{base} O\bot$ then $\vdash A \rightarrow OA$, and (ii) if $O(\neg A \& OA) \vdash_{base} O\bot$ then $\vdash OA \rightarrow A$.[49]
>
> **Proof:** (i) Suppose $O(A \& \neg OA) \vdash_{base} O\bot$. By the characterization of $\vdash$, $A \& \neg OA \vdash \bot$. By Reductio ad Absurdum, $A \vdash \neg\neg OA$. By Double Negation Elimination and Transitivity, $A \vdash OA$. By Conditional Proof, $\vdash A \rightarrow OA$. (ii) Suppose $O(\neg A \& OA) \vdash_{base} O\bot$. By the

[48] The former principle is related to what is sometimes called the "omissive version" of Moore's paradox ("A and I don't believe that A"). The latter principle is related to what is sometimes called the "commissive version" of the paradox ("A and I believe that it is not the case that A").

[49] This result suggests a natural technical question: Suppose we have a normal modal logic with modal operator O. What are the frame conditions corresponding to the modal principles $O(A \& \neg OA) \rightarrow O\bot$ and $O(\neg A \& OA) \rightarrow O\bot$? Using standard techniques, it is easy to show that the first principle is valid on all and only the frames that satisfy the following first-order condition: For every world w, if there is a world accessible from w then some world accessible from w is such that all worlds accessible from it are accessible from w. The second principle is valid on all and only the frames that satisfy the following first-order condition: For every world w, if there is a world accessible from w then some world accessible from w is such that some world accessible from it is accessible from w.

characterization of $\vdash$, $\neg A\&OA \vdash \bot$. By Reductio ad Absurdum, $OA \vdash \neg\neg A$. By Double Negation Elimination and Transitivity, $OA \vdash A$. By Conditional Proof, $\vdash OA\rightarrow A$.

What these results tell us is the following: Given Conditional Proof, even restricted to "ordinary" claims, we had better not try to characterize validity as the preservation of an operator O that either (i) obeys $OA \vdash_{base} OOA$ and is strong in the sense that $A\rightarrow OA$ is not always a validity, or (ii) obeys $OOA \vdash_{base} OA$ and is weak in the sense $OA\rightarrow A$ is not always a validity. Given Reductio ad Absurdum, Conditional Proof, Double Negation Elimination, and Transitivity, even restricted to "ordinary" claims, we had better not try to characterize validity as the preservation of an operator O that either (i) obeys $O(A\&\neg OA) \vdash_{base} O\bot$ and is strong in the sense that $A\rightarrow OA$ is not always a validity, or (ii) obeys $O(\neg A\&OA) \vdash_{base} O\bot$ and is weak in the sense that $OA\rightarrow A$ is not always a validity.[50] And even if we don't have Conditional Proof, we will still be able to derive $OA \vdash A$ and $A \vdash OA$ (respectively), which are bad enough.

Should we endorse a level-bridging principle, such as $\neg O(A\&\neg OA)$ or $\neg O(\neg A\&OA)$? This will presumably depend on the identity of the operator O. For familiar epistemic operators, such as *has justification for believing*, it is currently much debated whether such principles should be endorsed.[51] (The debate is typically focused on the question of whether epistemic akrasia – having a mismatch between one's beliefs and what one believes would be rational for one to believe – can be rational.) Determining whether we should accept such level-bridging principles is beyond the scope of this paper. But it is worth pointing out that such principles have at least some intuitive plausibility and are supported by at least some intuitive cases. They are also supported by more theoretical arguments. For instance, it could be argued that it would be unfair

[50] Notice that this result does not rely on Cut, only Transitivity. It also does not rely on Contraction.
[51] For a sample of recent discussions of the topic, see Christensen (2016), Horowitz (2014), Huemer (2011), Lasonen-Aarnio (2020), Littlejohn (2018), Smithies (2012), Titelbaum (2015), and Worsnip (2018).

(in some sense) for one to be held to epistemic standards that one is not in a position to be justified in believing, so there can never be a mismatch between what one is justified in believing and what one is justified in believing about what one is justified in believing. Alternatively, it could be argued that there is a tension between having a belief and believing that one is not justified in believing it, so there can never be a mismatch between what one is justified in believing and what one is justified in believing about what one is justified in believing. Perhaps these lines of thought are incorrect. But they have some intuitive pull, and so such level-bridging principles are worthy of serious consideration.

6. Rational Credence

There is a final kind of epistemic account of validity that I would like to discuss. This is the proposal to characterize validity in terms of rational credence.

On a natural way to develop this account, we first characterize a set of admissible credence functions, $\boldsymbol{AC}$. Then we say that a single-premise inference is valid just in case, on all admissible credence functions, the credence of the premise is less than or equal to the credence of the conclusion.[52] In symbols:

$$A \vdash C \text{ just in case } \forall Cr \in \boldsymbol{AC}\ Cr(A) \leq Cr(C)$$

[52] See Field (1977; 2015) for accounts of logical validity in this ballpark, inspired by technical results due to Popper (1959, appendices iv and v). Field (1977) defines a notion of probabilistic validity and argues "that the inferences licensed by classical logic are precisely those which are probabilistically valid" (p. 388). Field (2015) extends the account to a wide range of non-classical logics but now views the probabilistic condition as something like a consequence of logical validity rather than providing necessary and sufficient conditions for logical validity. More precisely, Field argues that "[t]o regard an inference or argument as [logically] valid is (in large part anyway) to accept a constraint" on rational credences (p. 42). Field makes use of conditional credences rather than unconditional credences, but this difference won't matter in what follows.

A natural way to understand the set of admissible credence functions is as containing all the credence functions that an ideally rational agent could have. This includes the ideally rational ur-credence function(s) as well as any credence function that can be arrived at from an ideally rational ur-credence function by updating on a possible body of information.[53] On a popular view, such credence functions will be probabilistically coherent. They will presumably also satisfy various other constraints, perhaps including having a maximal credence in analytic or conceptual truths, obeying some version of Lewis's principal principle,[54] and so forth.

If one would like to provide a definition of validity, one would have to provide an account of admissibility that doesn't itself appeal to validity. On the standard axioms for probability theory, it is an axiom that all validities get maximal probability.[55] So either one would have to provide an account of when a claim is a validity that is antecedent to the definition of valid inference, or one would have to make use of a different axiomatization of probability theory.[56] If, instead, one is not attempting to provide a definition of validity but merely some kind of characterization of it, one could view the account as providing a constraint on validity (or on rational credences, or on the connection between the two).

A credal characterization of validity has both intuitive and theoretical appeal. But it also runs into versions of the problem facing the other accounts of validity that we have discussed.[57] By

[53] The view is thus compatible both with the uniqueness thesis – that there is a single ideally rational ur-credence function – and with permissivism – the view that there is more than one ideally rational ur-credence function.
[54] See Lewis (1980).
[55] More precisely, it is an axiom that all logical validities get maximal probability. But this presumably should be extended to all validities.
[56] Field (1977) takes the latter option and makes use of an axiomatization (essentially) due to Popper.
[57] One could view $Cr(A) \leq Cr(C)$ as a claim about the preservation of an epistemic status – for all real numbers $x \in [0,1]$, the property of having a credence greater than or equal to x is preserved. One could then directly apply the results from earlier in this paper about the preservation of non-factive properties. The idea would be to let "O_xA" stand for $Cr(A) \geq x$. If $O_x(A\&\neg O_xA)\rightarrow O_x\bot$ is a plausible constraint on rational credence functions for every x, then we would have a version of the problem for non-factive epistemic operators. And similarly for $O_x(\neg A\&O_xA)\rightarrow O_x\bot$. The trouble is that the resulting pair of principles are not very compelling, at least if it can be rational to be uncertain what the rational credence is in a claim.

way of illustration, here are two plausible-seeming level-bridging principles for rational credence functions:

(i) For sufficiently small r, if $Cr(A) > Cr(\neg A)$ then $Cr(Cr(A) \leq r) \leq Cr(Cr(A) > r)$

(ii) For sufficiently large s, if $Cr(A) < Cr(\neg A)$ then $Cr(Cr(A) \geq s) \leq Cr(Cr(A) < s)$

These principles are fairly weak.[58] The first principle says that if the rational credence in A is in fact above the rational credence in the negation of A, then the rational credence that the rational credence in A is less than or equal to some low bound (e.g., .1) is less than or equal the rational credence that it is above that bound. The second principle says that if the rational credence in A is in fact less than the rational credence in the negation of A, then the rational credence that the rational credence in A is greater than or equal to some high bound (e.g., .9) is less than or equal to the rational credence that it is less than that bound. These principles are one way to capture the thought that one's rational credence in a claim cannot be too far out of alignment with what one rationally takes one's rational credence in the claim to be.

Given the first principle, we can argue that for any sufficiently small r and any admissible credence function Cr, $Cr(A \,\&\, \neg Cr(A) > r) \leq Cr(\neg A \vee Cr(A) > r)$. The argument is as follows: First suppose that $Cr(A) \leq Cr(\neg A)$. Since the rational credence of a conjunction is less than or equal to that of a conjunct, $Cr(A \,\&\, \neg Cr(A) > r) \leq Cr(A)$. Since the rational credence of a disjunct is less than or equal to that of a disjunction, $Cr(\neg A) \leq Cr(\neg A \vee Cr(A) > r)$. So by the transitivity of $\leq$, $Cr(A \,\&\, \neg Cr(A) > r) \leq Cr(\neg A \vee Cr(A) > r)$. Now suppose instead that $Cr(A) > Cr(\neg A)$. Since the rational credence of a conjunction is less than or equal to that of a conjunct, Cr(A &

[58] In orthodox probability theory, where $Cr(A) + Cr(\neg A) = 1$, these principles are equivalent to (i′) for sufficiently small r, if $Cr(A) > .5$ then $Cr(Cr(A) \leq r) \leq .5$, and (ii′) for sufficiently large s, if $Cr(A) < .5$ then $Cr(Cr(A) \geq s) \leq .5$. But for some non-orthodox probability theories – in particular, ones in which $Cr(A) + Cr(\neg A)$ can be less than or equal to 1 – the principles in the text are preferable.

$\neg Cr(A) > r) \leq Cr(\neg Cr(A) > r)$, which is equal to $Cr(Cr(A) \leq r)$. By the first principle, $Cr(Cr(A) \leq r) \leq Cr(Cr(A) > r)$. Since the rational credence of a disjunction is greater than or equal to the rational credence of a disjunct, $Cr(Cr(A) > r) \leq Cr(\neg A \vee Cr(A) > r)$. So by the transitivity of $\leq$, $Cr(A \& \neg Cr(A) > r) \leq Cr(\neg A \vee Cr(A) > r)$. Either $Cr(A) \leq Cr(\neg A)$ or $Cr(A) > Cr(\neg A)$. So by Reasoning by Cases, $Cr(A \& \neg Cr(A) > r) \leq Cr(\neg A \vee Cr(A) > r)$.

By the characterization of validity, $A \& \neg Cr(A) > r \vdash \neg A \vee Cr(A) > r$. By Conjunction Elimination, Reasoning by Cases, Explosion, Cut, and Contraction, $A \& \neg Cr(A) > r \vdash \bot$. By Reductio ad Absurdum, Double Negation Elimination, Transitivity, and Conditional Proof, $\vdash A \rightarrow Cr(A) > r$. In other words, it is a validity that if A then the rational credence in A is not low.

Given the second principle, using an analogous argument, we can show that for any sufficiently large s and for any admissible credence function Cr, $Cr(\neg A \& Cr(A) \geq s) \leq Cr(A \vee \neg Cr(A) \geq s)$. By Conjunction Elimination, Reasoning by Cases, Explosion, Cut, and Contraction, $\neg A \& Cr(A) \geq s \vdash \bot$. By Reductio ad Absurdum, Double Negation Elimination, Transitivity, and Conditional Proof, $\vdash Cr(A) \geq s \rightarrow A$. In other words, it is a validity that if the rational credence in A is high, then A.

These results are uncomfortable. It should not be the case that the rational credence in every truth is not low and that every claim that one has a high rational credence in is true. If one has too little evidence – or misleading evidence – the rational credence in a truth may be very low. And the rational credence in a falsehood may be very high. So these claims should not turn out to be validities.

This illustration relies on level-bridging principles that, while reasonably natural, don't appear in the literature. Here is a second illustration of how a problem can emerge that makes use of level-bridging principles that have been independently defended. Consider the following two principles: (i) for any $r \in [0,1]$, $Cr(A \mid Cr(A) \leq r) \leq r$, and (ii) for any $s \in [0,1]$ $Cr(A \mid Cr(A) \geq s) \geq s$. The second principle is essentially the principle that Dorst (2020) calls "Trust".[59] In orthodox probability theory, it is easy to show that these principles are equivalent. If rational credences obey orthodox probability theory, (i) can be used to show that for any $r \leq .5$, $Cr(A \,\&\, \neg Cr(A) > r) \leq Cr(\neg A \vee Cr(A) > r)$ and (ii) can be used to show that for any $s \geq .5$, $Cr(\neg A \,\&\, Cr(A) > s) \leq Cr(A \vee \neg Cr(A) > s)$. The argument for the former claim is as follows: It follows from (i) that $Cr(A \,\&\, Cr(A) \leq r) \leq r \times Cr(Cr(A) \leq r)$. So $Cr(A \,\&\, \neg Cr(A) > r) = Cr(A \,\&\, Cr(A) \leq r) \leq r \times Cr(Cr(A) \leq r) \leq r \leq .5$. So $Cr(A \,\&\, \neg Cr(A) > r) \leq Cr(\neg(A \,\&\, \neg Cr(A) > r)) = Cr(\neg A \vee Cr(A) > r)$. The argument for the latter claim is analogous. Then, just as above, we can argue for $\vdash A \rightarrow Cr(A) > r$ and $\vdash Cr(A) \geq s \rightarrow A$, respectively. So we end up with the same problematic "validities".

What these results show is that there is a clash between the proposed credal characterization of validity, plausible-seeming level-bridging principles, natural logical principles, and natural principles governing credences. Perhaps we should reject one of the logical principles that generates the problem, such as Reasoning by Cases or Conditional Proof. Or perhaps we should reject one of the credal principles. (Presumably, if we modify logic we should modify probability theory, and vice versa.) Perhaps we should give up the level-bridging principles at issue in the argument as well as other such principles that can be used to generate the problem. Or perhaps we should give up the proposed account of validity.

[59] Dorst's (2020) statement of Trust makes use of conditional credences. Given the assumption that the set of admissible credence functions, ***AC***, is closed under updating, and the assumption that updating works by conditionalization, the principle in the text is equivalent.

7. Conclusion

Let's take stock. There are reasons that support defining – or characterizing – validity in epistemic terms. (The one I find most compelling is that an epistemic definition of validity can explain its epistemic significance.) Nevertheless, the project of characterizing validity in terms of the preservation of some epistemic status runs into difficulty. The difficulty is that we seem to be able to show that certain claims are validities which should not count as validities since they are not even true. The argument that these claims are validities depends on a few natural logical principles as well as a sufficiently strong level-bridging principle. (In the case of a factive epistemic status, the level-bridging principle is factivity itself.) Similarly, if we characterize validity in terms of rational credence – so that an inference is valid just in case the rational credence in the conclusion is at least as great as the rational credence in the premise – we run into the same kind of difficulty.

These difficulties do not definitively show that epistemic characterizations of validity are to be rejected. There is always the prospect of responding by rejecting a logical (or credal) principle. There is also the prospect of responding by rejecting the relevant level-bridging principles. But the principles at issue are attractive ones, and it is uncomfortable to reject them.

There is one more option worth mentioning. That is to move to a different kind of epistemic characterization of validity. Perhaps the most natural such alternative is to move to a wide scope characterization. For instance, we could say that $A \vdash C$ just in case $\vdash_{base} O(A \rightarrow C)$, for some appropriate epistemic operator O. In the case of knowability, the idea would be that a single-premise inference from A to C is valid just in case it is knowable that if A then C. Similarly, for an account in terms of rational credences, we could move to the account that $A \vdash C$ just in case $\forall Cr \in \boldsymbol{AC}\ Cr(A \rightarrow C)=1$. A wide-scope account of validity seems to avoid the problems developed in this paper. However, this kind of account faces a different kind of problem, at least if it is taken to provide a definition of validity. In particular, such a definition seems incompatible with inferentialism about the conditional. The reason is that we cannot use the notion of validity to explain the meaning of the indicative conditional if we need to appeal to an attitude towards an indicative conditional to explain the meaning of validity. That would seem to be problematically circular.

Perhaps instead of an attitude towards a conditional, we could instead appeal to a conditional attitude. For instance, we could say that $A \vdash C$ just in case $\vdash_{\text{base}} M(C|A)$, where "$M(C|A)$" is the claim that one rationally has a conditional commitment in C given A. Alternatively, we could say that $A \vdash C$ just in case $\forall Cr \in \boldsymbol{AC}\ Cr(C|A)=1$, where "$Cr(C|A)$" stands for the rational conditional credence in C given A. That, too, would seem to avoid the problems.[60]

A different kind of alternative epistemic characterization is to characterize validity in terms of reasoning under suppositions (as well as beliefs). One way to do this would be to say that $A \vdash C$ just in case under any (or no) suppositions, if it is rational to accept A under those suppositions then it is rational to accept C under those suppositions.[61] The reason this avoids the difficulties is that there are coherent hypothetical cases in which A but it is not rational to accept A as well as coherent hypothetical cases in which ¬A but it is rational to accept A.

So there are options that can be pursued. But what all this shows is that whether an epistemic characterization of validity can be made to work is a delicate matter. And even if it can be made

[60] There is a potential problem here, too. If the notion of validity is used in defining the relevant conditional attitude, then we again run into a worry about circularity. Indeed, the same kind of potential problem arises for any definition of validity in terms of an epistemic status or attitude – it would be circular to both define validity in terms of a status or attitude and define the status or attitude in terms of validity.

[61] This proposal is analogous to Restall's response (in an unpublished 2009 blog post available at https://consequently.org/news/2009/06/01/rumfitt_part_1/) to Rumfitt's (2008, p. 80, n. 2) objection to Restall's (2005) bilateralist account of the normative force of (one kind of) validity. Ripley (2017, p. 311, n. 7) also endorses Restall's view. Thanks to an anonymous referee for another journal for bringing this exchange to my attention.

to work, either the correct epistemic characterization doesn't involve the preservation of an epistemic status from the premise(s) to the conclusion or, if it does, then some natural logical, credal, or epistemic principle has got to go.[62]

References

Anderson, C. A. (1983) "The Paradox of the Knower", *The Journal of Philosophy* 80: 338–55.

Armour-Garb, B., Stoljar, D. & Woodbridge, J. (2022) "Deflationism About Truth", *The Stanford Encyclopedia of Philosophy* (Summer 2022 Edition), E. Zalta (ed.), available at https://plato.stanford.edu/archives/sum2022/entries/truth-deflationary/.

Beall, JC. (2009) *Spandrels of Truth*. Oxford: Oxford University Press.

Brandom, R. (1994) *Making It Explicit*. Cambridge, Mass.: Harvard University Press.

Brandom, R. (2008). *Between Saying and Doing: Towards an Analytic Pragmatism*. Oxford University Press.

Brandom, R. (2018), "From Logical Expressivism to Expressivist Logic: Sketch of a Program and Some Implementations", *Philosophical Issues* 28: 70-88.

Carroll, L. (1895). "What the Tortoise Said to Achilles", *Mind* 4: 278-280.

Christensen, D. (2016) "Disagreement, Drugs, etc.: From Accuracy to Akrasia", *Episteme* 13: 397–422.

Dorst, K. (2020) "Evidence: A Guide for the Uncertain", *Philosophy and Phenomenological Research* 100: 586–632.

Dummett, M. (1991) *The Logical Basis of Metaphysics*. Cambridge, Mass.: Harvard University Press.

Elga, A. (2013) "The Puzzle of the Unmarked Clock and the New Rational Reflection Principle", *Philosophical Studies* 164: 127–39.

Field, H. (1977) "Logic, Meaning, and Conceptual Role", *The Journal of Philosophy* 74: 379–409.

Field, H. (2006) "Truth and the Unprovability of Consistency", *Mind* 115: 567–605.

Field, H. (2008) *Saving Truth from Paradox*. Oxford: Oxford University Press.

Field, H. (2015) "What Is Logical Validity?". In C. Caret & O. Hjortland (eds.), *Foundations of Logical Consequence*. Oxford: Oxford University Press, pp. 33–70.

Fitch, F. (1963) "A Logical Analysis of Some Value Concepts", *The Journal of Symbolic Logic* 28: 135–42.

Gibbard, A. (2003) *Thinking How to Live*. Cambridge, Mass.: Harvard University Press.

Harman, G. (1986) *Change in View: Principles of Reasoning*. Cambridge, Mass.: The MIT Press.

Hawthorne, J. (2004) *Knowledge and Lotteries*. Oxford: Oxford University Press.

Horowitz, S. (2014) "Epistemic Akrasia", *Noûs*, 48: 718–44.

[62] Thanks to Riki Heck, Stephan Krämer, Stephan Leuenberger, Bruno Whittle, and two anonymous referees for helpful comments.

Huemer, M. (2011) “The Puzzle of Metacoherence”, *Philosophy and Phenomenological Research* 82: 1–21.
Incurvati, L. & Schlöder, J. (forthcoming) “Inferential Deflationism”, *The Philosophical Review*.
Incurvati, L. & Schlöder, J. (2017) “Weak Rejection”, *Australasian Journal of Philosophy* 95: 741–60.
Lasonen-Aarnio, M. (2020) “Enkrasia or Evidentialism? Learning to Love Mismatch”, *Philosophical Studies* 177:597–632.
Lewis, D. (1980) “A Subjectivist’s Guide to Objective Chance”, in R. Jeffrey (ed.), *Studies in Inductive Logic and Probability*, Volume II. Berkeley: University of California Press, pp. 263–93.
Littlejohn, C. (2018) “Stop Making Sense? On a Puzzle About Rationality”, *Philosophy and Phenomenological Research* 96: 257–72.
Mackie, J. L. (1980) Truth and Knowability”, *Analysis* 40: 90–92.
MacFarlane, J. (2004) “In What Sense (If Any) is Logic Normative for Thought?”, presented at the 2004 Central Division APA symposium on the normativity of logic.
Makinson, D. (1965) “Paradox of the Preface”, *Analysis* 25: 205–7.
McGlynn, A. (2014) *Knowledge First?* New York: Palgrave Macmillan.
Moore, G. E., (1903) *Principia Ethica*. Cambridge: Cambridge University Press.
Moore, G. E. (1942) “A Reply to My Critics”, in P. Schilpp (ed.), *The Philosophy of G. E. Moore*. Evanston, IL: Northwestern University Press, pp. 533-677.
Murzi, J. & Shapiro, L. (2015) “Validity and Truth-Preservation”, in T. Achourioti, H. Galinon, J. Fernández, & K. Fujimoto (eds.), *Unifying the Philosophy of Truth*. Dordrecht: Springer, pp. 431-59
Paseau, A. (2008) “Fitch’s Argument and Typing Knowledge”, *Notre Dame Journal of Formal Logic* 49: 153–76.
Peacocke, C. (1992) *A Study of Concepts*. Cambridge, Mass.: The MIT Press.
Popper, K. (1959) *The Logic of Scientific Discovery*. New York: Basic Books.
Prawitz, D. (2012) “The Epistemic Significance of Valid Inference”, *Synthese* 187: 887–98.
Prawitz, D. (2015) “Explaining Deductive Inference”, in H. Wansing (ed.), *Dag Prawitz on Proofs and Meaning*. Springer, pp. 65–100.
Restall, G. (2005) “Multiple Conclusions”, in P. Hajek, L. Valdes-Villanueva, & D. Westerstahl (eds.), *Logic, Methodology and Philosophy of Science: Proceedings of the Twelfth International Congress*. Kings' College Publications, pp. 189–205.
Ripley, David. (2017) “Bilateralism, Coherence, Warrant”, in F. Moltmann and M. Textor (eds.), *Act-Based Conceptions of Propositional Content*. Oxford: Oxford University Press, pp. 307-24.
Rumfitt, I. (2008) “Knowledge by Deduction”, *Grazer Philosophische Studien* 77: 61-84.
San, W. K. (2020) “Fitch’s Paradox and Level-Bridging Principles”, *The Journal of Philosophy* 117: 5–29.
Smithies, D. (2012) “Moore’s Paradox and the Accessibility of Justification”, *Philosophy and Phenomenological Research* 85: 273–300.
Stalnaker, R. (1975) “Indicative Conditionals”, *Philosophia* 5: 269–286.
Steinberger, F. (2022) “The Normative Status of Logic”, *The Stanford Encyclopedia of Philosophy* (Winter 2022 Edition), E. Zalta & U. Nodelman (eds.), available at: https://plato.stanford.edu/archives/win2022/entries/logic-normative/.

Steinberger, F. and Murzi, J. (2017) "Inferentialism", in B. Hale, C. Wright, & A. Miller (eds.), *A Companion to the Philosophy of Language*, second edition. Malden: Wiley Blackwell, pp. 197–224.

Sundholm, G. (2012) "'Inference Versus Consequence' Revisited: Inference, Consequence, Conditional, Implication", *Synthese* 187: 943–56.

Tennant, Neil. (1987) *Anti-Realism and Logic*. Oxford: Oxford University Press.

Tennant, Neil. (1997) *The Taming of the True*. Oxford: Oxford University Press.

Titelbaum, M. (2015) "Rationality's Fixed Point (Or: In Defense of Right Reason)", *Oxford Studies in Epistemology* 5: 253–94.

Williamson, T. (2000) *Knowledge and its Limits*. Oxford: Oxford University Press.

Worsnip, A. (2018) "The Conflict of Evidence and Coherence", *Philosophy and Phenomenological Research* 96: 3–44.