Modest Inferentialism for Classical Logic*

Draft: Cite as such

Jack Woods

December 24, 2015

Modest Inferentialism, the view that inferential role *fixes* the meaning of (at least logical) expressions, against a background grasp of meaning, has had a bit of a checkered history. In this piece, I discuss problems with this view and criteria for a successful version of it. I show how the best extant version of modest inferentialism, due to James Garson, has trouble with these criteria and discuss what is needed to overcome these problems. To solve them, I develop an interpretation of the formal model-theoretic conditions which Garson-style modest inferentialism generates for the classical rules for the connectives which, in turn, motivates a principled restriction on admissible models. This involves a discussion of how to represent contingency in a modest inferentialist setting, the role of an intuitive interpretation in justifying side-conditions on models, and what atomic sentences represent in this setting. This interpretation satisfies the intuitive criteria for a successful modest-inferentialist account of the meaning of the logical connectives; it is a strong contender, I reckon, for an internally satisfying account of the meaning of the logical connectives—and one which does not extend to intuitionistic logic. This last point furnishes a not entirely disreputable argument against intuitionistic logic in favor of something approximating classical.

Introduction

Modest inferentialism holds that inferential practice or inferential role *fixes* the meaning of words against a background semantics. Of particular interest are *logical expressions* such as propositional connectives, quantifiers, and =. It is prima facie plausible that their meanings are fixed by their characteristic rules of implication; fixed in the sense of being specified by these rules against an

*Thanks to Catharine Diehl for comments on an earlier draft, to Patrick Blackburn and David McCarty for comments on a version delivered at ECAP 2014, and to Gil Sagi for a set of very useful comments. Thanks especially to James Garson for extended discussion of these issues and help with a number of issues in the paper. All mistakes which remain are my fault.

assumed account of the *type* of meaning such expressions should have.[1]

Modest inferentialism differs both from a fully semantic account of these expressions, where the rules of implication are held hostage to the meaning and from a pure inferentialist view where there is no sense in which the *meaning* of these expressions is to be understood in model-theoretic terms.[2] In contrast to these views, according to modest inferentialism we know more or less what kind of meaning a propositional connective # should have—say, a truth-function—but it is the role of the characteristic rules for # to specify *which* truth-function, among the contenders, this is.

So modest inferentialism takes our inferential practices to *define*, in some sense of 'define', the meaning of the logical connective. It does not, however, assume that we could come to grasp the meaning of the connectives by means of their characteristic rules—such a program would fall victim to the criticism of Carnap mooted in Quine (1936). We need logic to get logic out of definitions, so we cannot be expected to start from definitions and somehow come to grasp the meaning of the logical connectives—not unless we knew them already. Since modest inferentialism aims to account for the meaning of all logical expressions in terms of their inferential role, it cannot countenance the thought that we could have come to our knowledge of their meaning by means of acceptance of rules. It must rather leave etiological questions concerning inferential role to the side and see what meanings, if any, are specified by the role such expressions actually play.[3] The sensible goal is an internally satisfying account of the meaning of the logical expressions—internal in that we are entitled to use logical resources to generate this account, satisfying in being able to do so. My goal in this paper is to provide such an account by giving a determinate account of our background conception of meaning and show how it solves various problems that arise for a modest inferentialist approach. We start by briefly examining an example of modest inferentialist analysis, expose some problems with various approaches, and lay down some criteria.

Criteria for a Successful Modest Inferentialist Account

Consider conjunction, the most frequently cited success cases for this approach.[4] We can specify the inferential role of the conjunction by means of three rules:

[1] I use 'implication' as opposed to 'inference' in deference to Harman's point that the characteristic rules governing a connective are rules about what implies what, not rules, except indirectly, about how to *infer*. Inference is an action, implication a relation. Harman's point does not by itself scotch an inferential-role account of the meanings of the connectives—see Harman's own preferred such account in Harman (1986). The lengthy dispute about inferential role and the role of inference in reasoning would distract from the argument I give below; if necessary, read 'implicational role' for 'inferential role'.

[2] The best rigorous development of this latter line of thought is due to Prawitz (1985). The best non-rigorous development of this line of thought is due Brandom (1998).

[3] For useful discussion on this and related issues about justification, see Schechter (2013).

[4] See, for example, Peacocke (2004) and Boghossian (1996).

$$(\wedge^{I})\ \frac{\Gamma \vdash \varphi \quad \Delta \vdash \psi}{\Gamma, \Delta \vdash \varphi \wedge \psi} \qquad (\wedge^{O_L})\ \frac{\Gamma \vdash \varphi \wedge \psi}{\Gamma \vdash \varphi} \qquad (\wedge^{O_R})\ \frac{\Gamma \vdash \varphi \wedge \psi}{\Gamma \vdash \varphi}$$

Plausibly, this exhausts the role of $\wedge$, at least when it occurs as the main connective of a sentence (though, obviously, specifying its interaction with $\neg$ takes more work). If we require that $\wedge^I$ and $\wedge^O$ preserve truth and presume a bivalent semantics, then this trio of rules also exhausts the meaning of 'and' as it specifies completely the intuitively correct truth-function. The two rules $\wedge^{O_{L/R}}$ tells us that anytime $\varphi \wedge \psi$ is true, so is φ and ψ, likewise, that if either φ or ψ is not true (and so is false), then neither is $\varphi \wedge \psi$. And $\wedge^I$ tells us that if φ and ψ are true, so too is $\varphi \wedge \psi$. Together, they nail down the unique truth-function which assigns the conjunction $\varphi \wedge \psi$ true if and only both conjuncts are true.

This is a singular case; Carnap (1959) shows that the simple assumption of truth-preservation over the generated *consequence relation*—i.e. the result of closing the structural rule of identity under the relevant rules of inference and the structural rule of weakening—pins down the intuitively correct truth-function for only conjunction and the trivial truth-functions; Humberstone (1996) shows that truth-preservation over the *rules* yields classical truth-functions for intuitionistically acceptable sets of rules like the implicational fragment of classical logic.[5] So on either of these methods, we either get far too little—what we mean by these connectives far outstrips what the inferential role pins down—or far too much—the inferential role corresponding to a weaker meaning nevertheless picks out the stronger meaning.[6]

These results motivate criteria for a successful modest inferentialist program. First, and most importantly, the meaning fixed by the inferential role should not outstrip the inferential role. Peirce's law

$$((\varphi \to \psi) \to \varphi) \to \varphi$$

is not derivable in the implicational fragment of classical logic; applying a method for specifying the meaning of a connective to the implicational fragment of classical logic should not yield the full classical meaning of $\to$ on which Peirce's law is a tautology. We can glean a criterion for a successful account here: reading $\vdash_{\mathfrak{R}}$ for the proof-relation of a set of rules $\mathfrak{R}$, $\#^{\mathfrak{R}}$ for the meaning determined—in whatever way—by $\mathfrak{R}$ for a connective $\#$, and $\models_{\mathfrak{R}}$ for the consequence relation generated by letting each connective $\#$ mean $\#^{\mathfrak{R}}$:

$$\text{(A) } \Gamma \models_{\mathfrak{R}} \varphi \text{ if and only if } \Gamma \vdash_{\mathfrak{R}} \varphi$$

This is an intuitive form of soundness and completeness—the proof-relation (here indicating something like inferential role) should not outstrip the consequence relation we can generate from the meaning determined (in whichever

[5] See also Garson (2010).

[6] I don't mean 'weaker' or 'stronger' here to be taken too seriously. I am merely thinking of intuitionistic propositional logic as a sublogic of classical logic. Read my remarks accordingly.

way) by the proof-relation. And the consequence relation should likewise not outstrip the meaning relation, as happens in the case of Peirce's law when we use truth-preservation over the rules as our account of how rules determine meaning.

It turns out that this criterion is not sufficient; we not only want the meanings extracted from the rules to correspond to what can be derived by means of the rules, we also want the meanings to look roughly like the meanings we expect. For example, we think that meanings of the connectives are *compositional*—given only the meaning of the components of a sentence together with the meaning of the logical expressions in it, a semantic value should be determined for the sentence; *categorical*—roughly, that the meaning assigned to a sentence in terms of the meaning of the components in it is *unique*; and *complete*—inferential role should determine a meaning for a connective which yields meaningful sentences anywhere we would ordinarily take such a sentence to be meaningful. The third, the second, and plausibly the first are features of our intuitive understanding of the logical constants; they should be features of the meaning determined by inferential role if modest inferentialism is correct. We can articulate this criterion of adequacy in informal terms:

(C) Features of our ordinary understanding of the meaning of the logical constants ought to be preserved by the meaning assigned to them by their inferential role.

The idea behind (C) is simple—any plausible semantic candidate for the meaning of the connectives, this meaning will integrate into an overall semantic picture. Our usual semantic picture is compositional, categorical (at least as a default), and complete; the meanings we derive from the connectives shouldn't disturb this. If the derived meanings of the connectives failed to be compositional, categorical, and complete, then incorporating the derived meanings into our overall semantic picture would mean rejecting our usual semantic picture, contra the assumption we started off with that we had a grasp of the sort of meanings the connectives ought to have. Satisfaction of (C) is thus important to the overall plausibility of the modest inferentialist picture.

Joint satisfaction of (A) and (C) would motivate taking modest inferentialism very seriously as a foundational semantic view. Accepting modest inferentialism would allow us to satisfy both our intuition that we understand the meaning of the connectives while also allowing us to treat the inferential role of the connectives as of paramount importance in understanding their meaning. These intuitions pull in opposite directions and at least one is unsatisfied on each of the proof-theoretic or model-theoretic account of connective meaning; a hybrid view would be very welcome.

There are also other advantages to the modest inferentialist framework. For example, it provides a solution to inferentialism's *tonk problem* in the sense that it provides a decisive reason to exclude adding tonk to an existing set of proof rules.[7] Assume (A) and that the target set of rules $\mathfrak{R}$ we are analyzing contains the tonk rules:

$$(Tonk^I)\ \frac{\Gamma \vdash \varphi}{\Gamma \vdash \varphi\text{-tonk-}\psi} \qquad\qquad (Tonk^O)\ \frac{\Gamma \vdash \varphi\text{-tonk-}\psi}{\Gamma \vdash \psi}$$

and the rules for conjunction given above. Then, on the assumption that $\vdash_{\mathfrak{R}}$ is sometimes satisfied (an assumption guaranteed by the structural rule of identity), given any two sentences φ and ψ, $\varphi \vdash_{\mathfrak{R}} \psi$. By (A), this means that $\varphi \vDash_{\mathfrak{R}} \psi$, deeply conflicting with our intuitive sense of the meaning of logical connectives like conjunction. The intuitive meaning of conjunction is not such that $\varphi \wedge \psi$ is logically equivalent with every other sentence. And, of course, the situation here generalizes. So, adding tonk into a language guarantees (again, assuming the structural rule of identity) a massive violation of (C), giving us a decisive reason to exclude adding it to any set of rules.[8]

So the promiscuous tonk rules should be avoided since they bend the meaning of our actual connectives completely out of shape with our intuitive understanding of what their meanings are supposed to be. This sort of argument against tonk seems to me to be on better footing than the ongoing research program of specifying a notion of *harmony* that allows us to treat tonk and similar connectives as pathological though, admittedly, my approach isn't obviously open to a pure inferentialist.[9] Likewise, since in classical logic connectives like negation are not conservative over base languages like the implicational fragment, this approach also seems better than banning non-conservative extensions to a language (as in the approach described in Belnap (1962)).[10]

How to get (formal) Meaning from Rules

Of course, we need an actual method for extracting meanings from a set of rules to make sense of the modest inferentialist position. Garson (2001, 2013) has shown us the correct way to proceed.[11] We will say that an argument $\Gamma \vdash \varphi$ is *derivable* from a set of rules $\mathfrak{R}$ when and only when it can be obtained in the usual way from instances the structural rule of identity $\varphi \vdash \varphi$ by using rules

[7] This solution is open to any view that satisfies (A) and (C). Of course, solutions to tonk can also be given by non-inferentialist views and a similar solutions to the below can be given for views which are inferentialist about some fragment of a language (say, additions to a fixed language), but presume we already have fixed the meaning of another fragment. See Belnap (1962) for a general recipe for the latter view.

[8] This argument would not work if we actually spoke a language which contained tonk; it just tells us why we should not add tonk into a language we do speak. See Warren (2015) for a discussion of what a tonkish language would have to look like.

[9] For a discussion of recent work in this tradition, see Hjortland (2010). For throwing-the-baby-out-with-the-bathwater objections to this approach, see Warren (2015).

[10] Belnap's approach, however, is available to the modest inferentialist.

[11] Apart from typographical details, formulations, and small matters of exposition, the formal details and method in this section are due to Garson (2013) except where otherwise noted.

in $\mathfrak{R}$ and the structural rules of weakening and cut.[12] Given a valuation—here just an assignment of $\mathbf{T}$ and $\mathbf{F}$ to sentences of a language $\mathfrak{L}$—we'll say that an argument like $\Gamma \vdash \varphi$ *holds* in it when φ is $\mathbf{T}$ if every sentence in Γ is $\mathbf{T}$. Since we are not presuming that the meanings of the connectives will be truth-functional, we need to work with a notion of a model which is more general. So we will use not valuations, but sets of valuations (which I call "*general models*" in Woods (2012)) as our semantic framework. This allows *intuitionistic* truth conditions, versions of multi-valued logics, and the like to be contenders for the meanings of the connectives.[13]

Given a general model V, we will say that an argument is *correct* on V just in case it holds throughout V (i.e. holds in every $v \in V$). An arguments-to-argument rule R—indicated $\langle P_0, ..., P_n, C\rangle$—*preserves correctness* on V just in case the conclusion argument holds throughout V if the premise arguments hold throughout it. For example, $\wedge^{O_L}$—$\langle \Gamma \vdash \varphi \wedge \psi, \Gamma \vdash \varphi\rangle$—preserves correctness on any general model V where a conjunction is assigned $\mathbf{T}$ in every $v \in V$ when every sentence in Γ is just in case each conjunct is assigned $\mathbf{T}$ by every $v \in V$ when every sentence in Γ is.

Let $\mathfrak{V}_{\mathfrak{R}}$ be the set of general models on which every rule in $\mathfrak{R}$ preserves correctness. Then we immediately satisfy (A) since Humberstone (1996) has shown that:

1. A rule $\langle P_0, \ldots, P_n, C\rangle$ preserves correctness on every V in $\mathfrak{V}_{\mathfrak{R}}$ if and only if we can derive C from $P_0, \ldots P_n$ using $\mathfrak{R}$.

2. An argument $\Gamma \vdash \varphi$ is correct on V for every $V \in \mathfrak{V}_{\mathfrak{R}}$ if and only if $\Gamma \vdash \varphi$ is derivable from the rules in $\mathfrak{R}$.[14]

2 is a form of soundness and completeness result connecting the notion of correctness throughout a set of general models $\mathfrak{V}_{\mathfrak{R}}$ with a set of arguments-to-argument rules $\mathfrak{R}$. 1 guarantees that all *admissible* rules—rules which preserve correctness throughout $\mathfrak{V}_{\mathfrak{R}}$—for a connective are derivable. Given these results, the set of general models on which a set of arguments-to-argument rules preserves correctness very plausibly amounts to a perfect characterization of a sort of meaning generated by a set of rules. After all, these general models are all the ways of interpreting the sentences of the language where the rules aren't violated. We obtain a more familiar notion of meaning when we can characterize $\mathfrak{V}_{\mathfrak{R}}$ by means of a recursive condition.

For example, the meaning generated by the set of rules consisting of *modus ponens* and *conditional proof* ($\mathfrak{R}_{\rightarrow}$) is exactly what it ought to be—a class of

[12] See any reasonable textbook on natural deduction for details.

[13] See Garson (2001) for a defense of using this notion for our characterization of the meaning of the connectives. Since my purpose here is to show how to use Garson's method to give an internal justification of classical logic, I won't concern myself with accommodating more unreasonable logics such as those throwing out the structural rules. See Woods (manuscript) for reasons.

[14] See (Humberstone, 1996, 457-458) for more details and proof.

models which can be characterized by the *intuitionistic* truth-condition for $\to$. To see this, define a relation $\leq$ on V as follows. Let $v \leq v'$ iff $\forall \varphi\ v(\varphi) = \mathbf{T} \Rightarrow v'(\varphi) = \mathbf{T}$. A set of valuations V ordered under $\leq$ is a familiar structure, resembling a Kripke model for intuitionistic logic.[15] Then $(\mathfrak{R}_{\to})$ is characterized by means of a condition $C^{\to}$ as follows:

$$(C^{\to})\ \forall v \in V\ [v(\varphi \to \psi) = \mathbf{T} \text{ iff } \forall v' \in V,\ [v \leq v' \ \&\ v'(\varphi) = \mathbf{T}] \Rightarrow v'(\psi) = \mathbf{T}]^{16}$$

Since these two rules give the implicational fragment of both the classical and intuitionistic propositional calculus and since the intuitionistic conditional is weaker—in fact the weakest meaning which can be assigned to $\to$ in the presence of the structural rules and the deduction theorem—this is exactly the result we want. Similar results hold for & and $\neg$.[17]

However, as was shown in Woods (2012), the meaning this method assigns to disjunction is non-standard in a pair of ways. The meaning condition for disjunction is that $v(\varphi \lor \psi) = \mathbf{T}$ just in case:

$$\forall \chi\ v(\chi) = \mathbf{F} \Rightarrow \exists v'([v \leq v' \land v'(\chi) = \mathbf{F}] \land [v'(\varphi) = \mathbf{T} \lor v'(\psi) = \mathbf{T}])$$

But, first, this condition fails to generate a unique assignment of meaning to a disjunction on the basis of the meaning assigned to the disjuncts. Second, it fails to be compositional in the sense that the meaning of a disjunction can be fixed in some cases by fixing the meaning of a propositional variable which does not occur as a subformula of the disjunction Woods (2012). These both constitute significant violations of condition (C) argued for above. I do not think, though, that this scotches the program of giving an account of what meaning is fixed by rules by means of correctness preservation; far from it. Correctness preservation is the best account of the meaning conferred on logical expressions by means of their rules of implication—if this problem is devastating, it is modest inferentialism that is scotched.[18]

Is this problem devastating for modest inferentialism? As Garson (2013) demonstrates in some detail, we can lay down a side condition on the general models which excludes the troublesome cases Woods developed. However, Woods sketched a worry for this approach in his earlier work. The problem, in short, is that such conditions can and should be viewed as part of the rules

[15] See Kripke (1965).

[16] The formulation of $C^{\to}$ is mine.

[17] See Garson (2013) for details. The approach offered here and instigated by Garson and Humberstone bears a certain resemblance to later the possiblity semantics of Kit Fine and Ian Rumfitt. See, for example, Fine (2014) and (Rumfitt, 2015, Pt. 2). However, as their motivations are, at least largely, semantical instead of inferentialist, I will leave a detailed comparison to later work. Worries about the motivation of side-conditions like those discussed below reemerge on their pictures, though one has to look closely to detect where such assumptions are coming into play.

[18] Alternatively, you might think that the rules for $\lor$ are simply tonkish (Garson, 2013, pg. 273). I think this cure worse than the disease.

which generate the meaning of the connectives. But these side conditions violate intuitive criteria for how a rule should lay down an account of meaning. That is, they embody a violation of condition (C). This worry can be avoided, but to do so we need to think about how interpret the formal meaning conditions generated by Garson's method. And, of course, this is an important issue all by itself. A formal meaning condition like that given by $C^{\rightarrow}$ isn't yet an account of *meaning*—at best it is a constraint on what meanings we can assign to a connective. The formal structure against which we define the condition characterizing a connective itself needs to be interpreted.[19]

How to get (actual) Meaning from Rules—and a problem

As mentioned above, we need to interpret the formal structure we have defined our conditions against. **T** and **F** are easy enough at first glance; we can take them to be truth and falsehood or something similar. We can also interpret atomic sentences in terms of basic chunks of information—claims like "grass is green", perhaps. But there are other aspects of the structure we need to interpret. Note, for instance, that $C^{\rightarrow}$ expresses the meaning of the conditional in terms of the meaning of its component antecedent and consequent *and* the relation $\leq$. We thus need to not only understand the meaning of atomic sentences to understand this condition, we need also understand $\leq$. We need to have an interpretation of $\leq$ that makes sense in the context of how we understand the connectives (i.e. in the context of condition (C)) in order to claim that $C^{\rightarrow}$ delivers an account of the meaning of $\rightarrow$.

For some assumptions about background meaning, the interpretation is fairly straightforward. For example, the intuitionistically inclined can simply apply the standard interpretation of Kripke-models for intuitionistic logic, which immediately yields a recognizable interpretation of ($C^{\rightarrow}$). So interpreted, a conditional claim is true *at an information stage* (interpreting the valuations), if and only if any information stage *containing* (interpreting $\leq$) containing enough information to verify the antecedent also contains enough information to verify the consequent. Note, however, that the problem described above looks very difficult for the intuitionistic interpretation of general models; there is no reasonable way to lay down a intuitionistically acceptable side-condition which excludes my troublesome models Woods (2012). For others, it is significantly more complicated. The classical logician, for example, has a difficult time in interpreting conditions like $C^{\rightarrow}$ in a way which does due justice to the intuitive meaning of classical connectives.

There *is* a classical interpretation which works. In the remainder of the paper, I will spell this interpretation out. We start by interpreting each valuation as, more or less, a partial specification of a possible world and $v \leq u$ as indicating that u is a fuller specification than v. We then get a natural inter-

[19]See Burgess (2008) for the point about formal semantics in contrast to meaning.

pretation of the conditional. How does this work? By taking those sentences which the valuation assigns **T** to correspond to the information carried by the specification. The sentences a valuation says are **T** are, that is, those which indicate what the specification has so far specified. Since $u \geq v$ indicates that u is a "fleshing out" of the specification v, we require the models to be persistent (that is, that anything true at v is still true at v'. **F** is not falsity on this interpretation. It indicates "false or unsettled". Falsity at a valuation v is indicated by the condition that something is **F** at a valuation and at every valuation $u \geq v$.[20] If a sentence is not **T** at a valuation and it's not the case that it is **F** at that valuation and every extension thereof, we'll say it's unsettled. We have now interpreted all the machinery in the background formal notion of a general model. It motivates some restrictions on the model—for example, we can plausibly throw out the trivial valuation which specifies everything on grounds that there is no possible world where every sentence—I did not say every *atomic sentence*—is true.[21]

So interpreted, $\varphi \to \psi$ says something along the lines of "any specification of a possible world sufficient to specify φ as true also specifies ψ as true." When we add the classical negation, conjunction, and disjunction rules, and one additional side condition on models, we get a close cousin of classical semantics. We get a class of general models which validates all the classical rules of inference, but which is slightly weaker than classical semantics—for example, a disjunction $\varphi \vee \psi$ can be **T** at a valuation in a general model even if neither disjunct is **T** at that valuation. These possibility semantics—first formulated in Humberstone (1981)—do a good job at capturing the meaning of the classical connectives which satisfies (A) above. They are not quite the classical semantic theory, but as is well known, classical semantics is stronger than the meaning generated from the classical implication rules. Consider, for another example, that absent a determinacy operator, supervaluational semantics validates the classical proof rules. We thus should not expect to recover full classical meaning from the rules for classical logic. Something close enough will do.[22] But, even to get something close enough, we need the presence of a side-condition enforced on general models.

$$\text{(LF) } v(\varphi) = \mathbf{F} \Rightarrow \exists v'[v \leq v' \ \wedge \forall v''(v' \leq v'' \Rightarrow v''(\varphi) = \mathbf{F})]$$

What this condition says is that any sentence **F** at a model is **F** at some

[20] See (Garson, 2013, 8.3) for a definition of "unsettled" in terms of negation. Since we are not working in this context with negation assumed, we let "unsettled" be neither true nor false.

[21] Though we need not do this. See Humberstone (1996) for discussion.

[22] The small details of the semantic meanings are not important for my account. But, just to be explicit, they are the combination of the intuitionistic meanings for the connectives with the side-condition (LF). The truth-conditions characterizing these meanings are very natural; for example, a disjunction is true at a valuation v if, at every extension of v, there is an extension u of that such that one or other of the disjuncts is **T** at u. Roughly, a disjunction is true at a valuation if there is a valuation extending it *from which* one or the other disjuncts will eventually be true whichever way we proceed.

extension and all *its* extensions. That is, that any sentence not yet settled at a specification is eventually settled false at some further specification.[23] We can also formulate a version, (pLF), defined only on atomic sentences instead of the entire language.

$$\text{(pLF) } v(p) = \mathbf{F} \Rightarrow \exists v'[v \leq v' \ \wedge \forall v''(v' \leq v'' \Rightarrow v''(p) = \mathbf{F})]$$

The question we need to ask is whether we can make sense of accepting (LF) on the basis of our interpretation. That is, whether (LF) intuitively follows from our understanding of the valuations as partial specifications of possible worlds and $\leq$ as indicating the "fleshing out" or "filling in" relation. To answer this, we need to note a few things. First, (LF) entails (pLF), but these conditions are *not* equivalent. Consider, for example, the following atomic general model defined over three propositional variables ($w \longleftarrow u$ indicate $w \leq u$):

$$\begin{tikzcd} v\ (p,r) \arrow[dr] & & u\ (q,r) \arrow[dl] \\ & w\ (r) & \end{tikzcd}$$

(pLF) holds in it. (LF) too, trivially, since we have no logically complex formulas. But consider adding the rules for $\vee$. We can then flesh out the general model so that

$$\begin{tikzcd} v\ (p,r,p \vee q) \arrow[dr] & & u\ (q,r,p \vee q) \arrow[dl] \\ & w\ (r) & \end{tikzcd}$$

where $p \vee q$ *does not* hold at w. (LF) thus fails in the general model while (pLF) holds.[24]

Second, (pLF) and (LF) are *side conditions* on general models—that is, a condition on acceptable general models that does not correspond to any meaning derived from the rules. Its role is to cut down the class of acceptable general models.[25] In our present context, it corresponds to a part of our intuitive understanding of the meaning to be generated by the rules.

[23]This, again, corresponds to Garson's notion of being *settled false* at a valuation v. See also Humberstone (1996).

[24]For a proof that the latter model is admissible given $\mathfrak{R}_\vee$, see Woods (2012).

[25]Garson and I disagree on the role of side-conditions—for him, a side condition is merely something which cannot be cast in truth-conditional terms (personal correspondence). This seems to me to be too weak of a condition for side conditions, given the sort of examples which motivate it and, in particular, our *antecedent grasp* on the meaning we are trying to model as the example of tense logic below suggests. Otherwise we would need accept that the rules for a connective express things which are not part of the meaning. Far better to allow that part of the meaning of a connective cannot be cast in pure truth-conditional form.

Garson gives the illustrative example of density for tense logic. If we demand that the models of tense logic be dense then we can recover $F\varphi \to FF\varphi$—anything which is going to be true is going to be going to be true—from the rules for F. Intuitively, this does not modify the meaning of F. It is just that the meaning of F, along with a particular class of designated general models, nails down as valid a schema of tense logic which is not derivative from the meaning of F itself. We are just interpreting the class of models over which we impose our constraints a particular way. If our aim is to model a way time could be—in particular, if our aim is to model what tense operators would mean on a dense interpretation of time—then we need not think that the density of time has anything to do with the meaning of F. It is rather that, given the density of time, we get a particular interpretation of the tense operators.

Of course, in the presence of density side condition, we cannot be said to be giving a general account of "it will be the case at some time that", but only an account of this in the context of a dense ordering on times. For an account of the general meaning of "it will be the case at some time that", we need to have no such restrictions on our general models. A similar caveat seems reasonable for the question of the meaning of the logical connectives.[26] In order to get an account of their meaning, we need a general interpretation that makes no substantial assumptions about the contents they are applied to. Something like this seems to underly the common criteria that logic should be *general* and *topic-neutral*—its characteristic expressions should be meaningful over a wide variety of possibilities. The possibility framework, minus (LF) or (pLF), has a *prima facie* claim to being general enough a framework to recover reasonable meanings of the logical connectives as it is a weakening of the idea of defining meaning over possible worlds. But what about the possibility framework with (pLF) or (LF)?

To say more about this, we need to be even more definite about how to interpret the general models. I will interpret them as representing partial specification of a possible world which are *epistemically* available to us.[27] Focusing on a particular valuation v, we view extensions of it as representing possibilities *consistent* with what we are currently hold as settled. What we hold as settled is represented by those sentences marked true or false at v. The fact that an atomic sentence is unsettled at a valuation means that it can go either way: to **T** where every further extension marks it as **T**, or **F** where every further extension marks it as **F**. This does not mean that a pair p, q of unsettled atomic sentences can both go true or false—this will depend on the content they are

[26] Tense logic is not logic in the relevant sense. Or, anyways, so I believe. I can't defend this view here, but see Harman (1972) for a general case for not treating tense or modal logic as logic and Williamson (forthcoming) for contrary considerations. Anyways, nothing substantial turns on this point.

[27] Clearly this epistemic interpretation is very lose and will require much idealization, but let that pass. I aim to just sketch the form of the solution here; details will come in separate work. Similar non-epistemic interpretations can obviously be given, but I won't pursue them here.

meant to represent. So how do (pLF) and (LF) fare on this interpretation?

Atomic Sentences, Contingency, and (pLF)[28]

An intuitive justification for (pLF), on the possibility framework, isn't too difficult to find—at least if we interpret partially specified possible worlds as representing something like epistemic situations (more on this below.) Atomic sentences are *typically*, though not always, intended to model contingent content. Such sentences should generally be capable of being both true and false in some sense. Further, they should be capable of being settled true or false if they are unsettled—we may, of course, let some some fragment of them be true or false throughout the model. Atomic sentences are meant to represent bits of content prior to interaction with the logical connectives; this does not mean that their only meaning is the truth-value we assign to them. Rather, we see the truth-values assigned to them in a valuation and, more generally, their profile across valuations in a general model, as representing the contours of the underlying content they stand for.[29]

I am not talking here about atomic sentences which represent content which is not intuitively contingent. It's often assumed that certain claims are assumed to be held true (or false) by any rational agent—say, mathematical claims, analytic truths, etc. We will presume that such sentences are represented in general models by being either **T** or **F** throughout the valuations in it, as appropriate to their content.[30] But, again, pairs of sentences might be jointly incompatible while each independent sentence is epistemically possible. So, for example, suppose we wanted to model a pair of logically consistent, but content-contrary statements like 'Bob is lost his game' and 'Bob won his game'. We could represent this in a general model by letting p represent the first, q the second, and having no valuation v in our general model where $v(p) = v(q) = \mathbf{T}$.

Remember that being marked **F** isn't quite being marked false; it's being either false or unsettled. Since we ought to be able to settle anything contingent and currently unsettled as definitely false in some fuller specification there should always be a way of extending a specification in order that something not yet specified becomes false—i.e. to a valuation where there is no fuller valuation where it is marked **T**. Note, importantly, that if a valuation has no extensions, then every sentence marked **F** is thereby false. So, to be more careful, any sentence marked **F** at a valuation which is not false at that valuation (i.e. there is an extension of that valuation where it is **T**) ought to be capable of being settled as false as well.

[28] This section is deeply indebted to an email back-and-forth with James Garson.

[29] See Woods (forthcoming) for a discussion of problems which arise if we do not pay attention to this fact about the interpretation of models.

[30] Not much turns on this point; I just raise it to fend off an obvious objection. Details would obviously have to specified here to give a full account.

To see the plausibility of this approach, consider an atomic general model where pLF fails:

$$\begin{tikzcd} v_0 & v_1 \arrow[l] & v_2 \arrow[l] & \dots \arrow[l] \\ v_0^*(p) \arrow[u] & v_1^*(p) \arrow[u] & v_2^*(p) \arrow[u] & \dots \end{tikzcd}$$

Both (LF) and (pLF) rule this out. But there is something strange about the model to begin with if we interpret general models as representing partial specifications of possible worlds. If p is contingent, as suggested, then why is it that we have no extension where it is definitively specified as false? Our general model seems incomplete.[31] Remember here that we are not, like the intuitionists, interpreting the valuations as representing moments of time or information stages. Rather we are interpreting them as things which are epistemically possible relative to a valuation v. If, relative to v_0, p is unsettled (as it is), then it seems that there should be a valuation extending v_0 where it is false since that is epistemically possible relative to v_0.

We can also give a direct argument for why we need (pLF) to hold on this interpretation. Suppose we want to intepret an atomic sentence p as unsettled at a valuation v. Since our relation $\leq$ is persistent, if at some extension u, $u(p) = \mathbf{T}$, this guarantees that at every $w \geq u$, $w(p) = \mathbf{T}$. So p really is settled as true at u. But, absent (pLF), we cannot say the same about settling as false since $\mathbf{F}$ isn't persistent. In order for a contingent unsettled sentence like p to be settled false, we need that there is an extension where it is $\mathbf{F}$ and where it *stays* $\mathbf{F}$. This is what (pLF) guarantees. And, therefore, (pLF) plays an essential role in our ability to interpret this formal framework in the way I've suggested. Moreover, this explains why we need no similar principle saying that if a sentence is unsettled, then it can settled as true.[32] This is captured immediately by the combination of the persistence of $\leq$ and the definition of being unsettled—if p is unsettled, then there is an extension where it is not $\mathbf{F}$. This means that there is an extension where it is $\mathbf{T}$, and persistence guarantees that it stays $\mathbf{T}$ at every extension thereof.

This sketchy justification for (pLF) seems on the right track. But what about justifying (LF) as a side-condition? Can we say the same thing about

[31] (pLF) does not require that a general model have *atomic maximal* valuations which have no extensions. Rather, it only requires that, for every atomic sentence which is $\mathbf{F}$ at some specification, there is some specification where it is false.

[32] Humberstone (1981) uses a condition, refinability, which says that if a sentence is unsettled, then it can be settled true and it can be settled false. This is unavailable to us here since our definition of $u \leq v$ is that everything $\mathbf{T}$ at u is $\mathbf{T}$ at v. If we required that every sentence $\mathbf{F}$ at a valuation become $\mathbf{T}$ at some extension of it, then we would never have a sentence which was actually false. Humberstone doesn't have a problem here because his $\leq$ varies somewhat independently of the valuations it is defined over (as long as it remains persistent.) Since we have found $\leq$ in our general models and not added it as a separate piece of apparatus, this route is not available to us. Note also that the definition of $\leq$ is crucial to the proofs that conditions like $C^{\vee}$ and $C^{\rightarrow}$ correctly characterize $\mathfrak{R}_{\vee}$ and $\mathfrak{R}_{\rightarrow}$. Thanks to James Garson for discussion of this very last point.

it? Not while we maintain that the meaning of logically complex formulas is derivative from just the rules and the background account of meaning. There does not seem to be anything natural to say about our grasp of the meaning of logically complex formulas which entails that any complex formula unsettled at a particular model must be settled false in some extension thereof. Consider our triangular example above. $p \vee q$ is $\mathbf{T}$ at every extension of w. But, as neither is true at w, given the meaning of $\vee$ forced by $R^{\vee}$, nothing forces us to say that $p \vee q$ is true at w. Why couldn't this model our actual situation. Both p and q are unsettled and once we settle either, we thereby settle $p \vee q$. But we are currently not required to force it by the rules for $\vee$ and we are not required to force it by our background understanding of the meaning of contingent content as represented by atomic formulas. So it seems doubtful that we can justify accepting (LF) as a side-condition on general models.

Can we accept stronger principles than (pLF) about atomic sentences? Suppose, for example, that we have a set Δ of some atomic sentences that are $\mathbf{F}$ at a valuation v. Should we assume that there is an extension where every sentence in Δ is settled as false? No, we should not. Just as a set of atomic sentences might represent contingent, but contrary content, a set of sentences might also represent contingent, but subcontrary content. That is, content such that if one of them is settled as false, the rest of them must be settled as true, but either could be settled as false from the standpoint of a valuation v where they are both currently unsettled. We do not want our models to exclude this possiblity, as they would if we accepted an analogue of (pLF) for sets of sentences. One the face of it, (pLF) is *exactly* what we need in order to make sense of our interpretation of the formal features of general models; moreover, this interpretation seems apt to the generality and universality features of logic that I mentioned above. But is (pLF) enough to vindicate our interpretation? Not yet; there is still a lingering problem.

The Remaining Problem and a Solution

(pLF) removes the compositionality problem for $\vee$. Taking (pLF) as a side-condition on acceptable general models—i.e. accepting that any admissible general models modeling our grasp of basic non-logical contents as represented by propositional variables will obey (pLF)—leaves two problems. First, $\vee$ is still not categorical—we can still assign $\mathbf{T}$ or $\mathbf{F}$ to disjunctions at valuations in certain models without disturbing the underlying atomic model. This still amounts to a violation of (C). Second, double negation elimination—$\langle \Gamma \vdash \neg\neg\varphi, \Gamma \vdash \varphi \rangle$—the rule that separates classical from intuitionistic negation, expresses the condition we have called (LF). That is, if we go on to accept full classical negation, then we close the gap between (pLF) and (LF) and thereby rule out both the categoricity and the completeness problem described above.

Unfortunately, this means that accepting double negation elimination not only modifies the meaning of $\to$, as we have long known from discussions of

Peirce's law, but it also *modifies* the meaning of $\vee$, in the sense that it eliminates the bizarre features of the meaning of $\vee$ described above. It eliminates certain atomic general models which is a very strange consequence in the present context. It is unintuitive that the meaning derived from the rules should enforce a constraint on the behavior of atomic sentences. This also seems like a violation of (C).[33] The first problem is serious if we require that each connective be *separable* in the sense that it is all and only the rules for that connective which specify its meaning. But if we reject this shackle, we can give a more cogent defense inspired by Garson's side condition suggestion. And, surprisingly, doing so also solves the second problem. This is exactly the strategy I want to pursue in the remainder of the paper.

Here is what I have in mind. Garson, in discussion what the classical rules for negation express, says

> While $\mathfrak{R}_{\sim}$ expresses the conjunction of $C^{\neg}$ with (pLF), the latter is not part of the truth conditions expressed by $\mathfrak{R}_{\sim}$. While $\mathfrak{R}_{\sim}$ brings with it failures of conservation and functionality, those failures are not to be blamed on the truth conditions expressed by $\mathfrak{R}_{\sim}$.[34]

But this is misleading. $\mathfrak{R}_{\sim}$ and $\mathfrak{R}_{\rightarrow}$ together imply Peirce's law, so any general model satisfying the conjunction of what they express ought to vindicate it. Consider

$$\begin{tikzcd} v\ (p) \arrow[dr] & & u\ (q) \arrow[dl] \\ & w & \end{tikzcd}$$

This can be fleshed out so as to respect the meaning of negation, the conditional, disjunction *and*, in line with what we have remarked above, in such a way that $w(p \vee q) = \mathbf{F}$. Note also that the resulting model satisfies (pLF), though not (LF).[35] Let φ be $p \vee q$. Consider now Peirce's law in the instance $((\varphi \to p) \to \varphi) \to \varphi$. $\varphi \to p$ is $\mathbf{T}$ only at v. So

$$v((\varphi \to p) \to \varphi) = u((\varphi \to p) \to \varphi) = w((\varphi \to p) \to \varphi) = \mathbf{T}$$

But, since $w \leq w$, $w(((\varphi \to p) \to \varphi) \to \varphi) = \mathbf{F}$ and the model does not validate Peirce's law. $C_{\neg}$ with (pLF) are not expressed by $\mathfrak{R}_{\sim}$, but rather $C_{\neg}$ with (LF). It is only with (LF) or equivalent principles that we get Peirce's law.

What Garson has in mind is that *in the context of a stronger meaning for* $\vee$—one where a disjunction is $\mathbf{T}$ at a valuation v if, at every $u \geq v$, at least one of the disjuncts is $\mathbf{T}$ at some extension of u—the conjunction of (pLF) and $C^{\neg}$

[33] See Garson (2013) and Woods (2012) for discussion of this problem.

[34] I have modified the quote to bring it in line with my favored notation.

[35] To see that it does not satisfy (LF), $p \vee q$ is true at both extensions of w. But every atomic sentence which is $\mathbf{F}$ at w is $\mathbf{T}$ at some extension.

is what is expressed by adding $\mathfrak{R}_\sim$ to the language. This is because in the context of this stronger meaning for disjunction, (pLF) and (LF) are equivalent. But since $\mathfrak{R}_\vee$—in fact, even $\mathfrak{R}_{\vee,\rightarrow,\neg,\wedge}$—do not express the stronger meaning for $\vee$, it is inappropriate to use the stronger meaning in determining what $\mathfrak{R}_\sim$ expresses. The only sense in which $\mathfrak{R}_\sim$ expresses the conjunction of (pLF) and $C^\neg$ is due to the fact that, making use of (LF) or the stronger meaning of $\vee$ which we obtain by means of (LF), (pLF) and (LF) are equivalent. But that is, again, because $\mathfrak{R}_\sim$ expresses something stronger than (pLF).

Without (LF) or the stronger meaning of $\vee$, we cannot recover double negation elimination or Peirce's law from $C^\neg$ and (pLF). The meaning expressed by $\mathfrak{R}_\vee$ is $C^\vee$. We need something like (LF) to get our general models into line. But this means that the failures of conservativeness *are* to be blamed on what is expressed by $\mathfrak{R}_\sim$ unless we can treat (LF) as a side condition motivated by our interpretation of the general models. But (LF) does *not* look like it captures anything basic about our understanding of non-logical content as embodied in our epistemic conception of possibilities, as discussed above, so it is implausible to view it as a side condition. So, summing up, the presence of classical negation disrupts the meaning of $\vee$ and $\rightarrow$ by expressing (LF), which in turn validates Peirce's law and nails $\vee$ down to a unique and categorical meaning.[36]

Instead of treating (LF) as a side condition, we should treat it as part of the meaning of classical negation, the part embodied in double negation elimination. We want to be able to represent "being unsettled", so we need to be able to say that for every propositional variable p and every valuation v, if p is unsettled at v, there is an extension where it is settled false and one where it is settled as true. This captures what it is for atomic sentences to represent content which is unsettled at a valuation.[37] In order for this account of being unsettled to go through, we need (pLF). (pLF), in turn, rules out the compositionality problem though not the categoricity problem. However, the full set of classical implication rules suffices to express (LF) since these rules contain double negation elimination. This rules out the categoricity problem, forcing every meaning-condition for a connective into a standard form where unique semantic values are assigned to each formula given the underlying value of its atomic subsentences. The resulting semantics easily satisfies (C) and, of course, continues to satisfy (A).

The essential feature of my solution is breaking the imposition of (LF) into

[36]Garson takes the notion of a side-condition to be cashed out purely in terms of being capable of being put into a (recursive) truth-condition for the connective and he notes that (LF) cannot be so treated. Note, though, (pLF) can be so treated:

$$v(p) = \mathbf{F} \text{ iff } \exists w \geq v\ w(p) = \mathbf{F} \wedge \forall u \geq w\ u(p) = \mathbf{F}$$

This strikes me as another reason to view (pLF) as capturing something about our treatment of non-logical content.

[37]Remember here that if there is no extension of v where it is $\mathbf{T}$, then it is *already* settled false at v. So if it is truly unsettled, then there is already an extension of v where it is $\mathbf{T}$.

two parts. (LF), simply imposed, is implausible. The only thing that should make a demand on acceptable atomic models should have to do with the meaning or interpretation of atomic sentences, such as our justification of (pLF) in terms of representing contingent unsettled content. (LF) also cannot be motivated by features of the rules (excluding double negation elimination) and features of our interpretation of general models. If, however, (pLF) is antecedently imposed as a side-condition on admissible general models, analogously to imposing density on a tense logic, then (LF) can then be interpreted, not as a side-condition, but rather as part of what is expressed by classical negation. It modifies the meaning of $\vee$ and $\to$, but this is just to say that the meaning of the classical set of inference rules is *holistic* in the sense that the full meaning of $\vee$ and $\to$ cannot be seen absent the full classical meaning of $\neg$.

In effect, we need and ought to read the truth-conditions for the connectives off of the rules *in concert*. (pLF) is simply too weak to remove all the problems with $\vee$ and $\to$. But (LF) is not, so if we can swallow the smidgen of holism that double negation elimination brings, there is no further problem. We obtain a satisfying account of the meaning of the classical connectives for our interpretation of general models which satisfies (A) and (C) and finds these meanings entirely in the classical rules for the connectives. Of course, this is not to say that no other interpretation is possible, but such an interpretation will have to face the problems I have laid out above and do so in a similar way. My goal here has been to articulate a reasonable interpretation of general models which (a) captures a notion of meaning that we are familiar with; (b) is uniform and general in the sense demanded by an account of the meaning of the *logical* connectives; and (c) satisfies (A) and (C) above. This I have done.

Conclusion

If we take the meaning of $\vee$ to be given solely by $C^{\vee}$ and the meaning of $\to$ as $C^{\to}$, then we cannot escape the fact that things which are not part of the meaning of the connectives, and not plausibly side conditions on general models, can change the set of true complex formulas involving that connective. In fact, a not entirely disrespectable argument for the primacy of classical inference rules as opposed to intuitionistic ones can be constructed out of this material. If we hew to a modest inferentialist framework, then we only have categoricity for $\vee$ if we have not only (pLF), but also (LF). But (LF) is simply too strong for an intuitionistic framework–it licenses intuitionistically unacceptable rules such as Peirce's Law. (pLF) itself is (a) too weak to solve the categoricity problem and (b) intuitively intuitionistically unacceptable Woods (2012).[38] But, with-

[38]It would distract from the central line of argument to go into this in detail, but it is worth noting that it is difficult for the intuitionist to explain the contingency of atomic sentences using features of general models (suitably interpreted). This is because it is difficult for the intuitionist to distinguish, using features of general models, between true contingency and merely never being settled as false. If we define "settled false" as "not **T** in any extension", contingency, understood as "not settled false and not settled true" amounts to "could be verified true in some extension, but not settled false in every extension". Depending on the

out (LF), $\vee$ does not have a categorical meaning. This means that we have a violation of (C) unless we go (quasi-)classical. But if, on the other hand, we are willing to accept holism about connective meaning, then we can motivate a satisfying account of the meaning of the classical connectives. Modest Inferentialist frameworks, in their strongest and most plausible incarnation, are classical.

Are there are other conditions like (pLF) which could be motivated on the basis of some understanding of the atomic sentences of general models? Perhaps, but on the interpretation I have suggested and nearby alternatives, none suggest themselves immediately. The burden is clearly on the classical modest inferentialist who wishes to resist my approach or the intuitionistic modest inferentialist to supply such a condition. One possible condition, suggested by an anonymous reviewer, is that we not rule out any distribution of **T** and **F** to atomic sentences. That is, that any possible distribution of **T** and **F** to the atomic sentences occur somewhere in a general model. But this requirement does not make sense on our interpretation—since we are considering atomic sentences to represent (a) bits of content which might, in fact, conflict prior to the action of the logical connectives and (b) the valuations in our model represent epistemic states of being sure of things; perhaps there are things which we *must* be sure about, such as the fact that there is something that exists. We can capture the reviewer's intuition by means of the fact that for any **F**-distribution to the contingent atomic sentences can be present as a valuation in some model or other. As far as logic is concerned, any **F**-distribution could serve as the set of things we are currently sure about of the things we could or could not be sure about.

To sum up, we can give an account of how to read meaning off of the rules for the connectives, given a particular formal set-up, and interpret this formal set-up in a natural way. The resulting account, given the classical set of rules, yields categorical, compositional, complete meanings which, though not *boolean*, are still recognizable as classical meanings. But to do this, we need to accept a side-condition, (pLF) on general models, which rules out certain atomic models. This condition can be motivated by interpreting it in terms of our understanding of the role of atomic sentences. (LF), which we need to solve the categoricity problem, can be viewed as extending the side-condition (pLF) on models by means of an inference rule, double negation elimination. Since the resulting account satisfies (A), it validates all and only those sentences which are classically valid. The account thus is capable of vindicating the intuition that the rules determine the meaning of the connectives as well as vindicating the idea that the meanings of the connectives are a familiar type of semantic condition—albeit modestly holistic ones.

---

particular interpretation of the general models on offer, this seems to conflate contingency with merely never being verified or falsified. Anyways, extended discussion would take us astray, but the intuitionist faces, at a minimum, a non-trivial hurdle here. Thanks to James Garson for discussion of this point and raising the initial worry about contingency for the intuitionist.

References

Belnap, N. D. (1962). Tonk, Plonk and Plink. *Analysis*, pages 130–134.

Boghossian, P. (1996). Analyticity reconsidered. *Noûs*, 30(3):360–391.

Brandom, R. (1998). *Making it Explicit: Reasoning, Representing, and Discursive Commitment*. Harvard University Press.

Burgess, J. (2008). Tarski's Tort. *Mathematics, Models, and Modality.*

Carnap, R. (1959). *Introduction to Semantics and Formalization of Logic*. Harvard University Press.

Fine, K. (2014). Truth-maker Semantics for Intuitionistic Logic. *Journal of Philosophical Logic*, 43(2-3):549–577.

Garson, J. (2001). Natural Semantics: Why Natural Deduction is Intuitionistic. *Theoria*, 67(2):114–139.

Garson, J. W. (2010). Expressive Power and Incompleteness of Propositional Logics. *Journal of philosophical logic*, 39(2):159–171.

Garson, J. W. (2013). *What Logics Mean: From Proof Theory to Model-theoretic Semantics*. Cambridge University Press.

Harman, G. (1972). Is Modal Logic Logic? *Philosophia*, 2(1):75–84.

Harman, G. (1986). The Meaning of Logical Constants. *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, pages 125–134.

Hjortland, O. (2010). The Structure of Logical Consequence: Proof-Theoretic Conceptions.

Humberstone, L. (1981). From Worlds to Possibilities. *Journal of Philosophical Logic*, 10(2):313–339.

Humberstone, L. (1996). Valuational Semantics of Rule Derivability. *Journal of Philosophical Logic*, 25(5):451–461.

Kripke, S. A. (1965). Semantical analysis of intuitionistic logic. In Crossley, J. and Dummett, M. A., editors, *Formal Systems and Recursive Functions*. 92–130. North-Holland, Amsterdam.

Peacocke, C. (2004). Understanding logical constants: A realist's account. *Studies in the Philosophy of Logic and Knowledge*, page 163.

Prawitz, D. (1985). Remarks on Some Approaches to the Concept of Logical Consequence. *Synthese*, 62(2):153–171.

Quine, W. (1936). Truth by Convention. *Philosophical Essays for Alfred North Whitehead*.

Rumfitt, I. (2015). *The Boundary Stones of Thought: An Essay in the Philosophy of Logic*. OUP Oxford.

Schechter, J. (2013). Could Evolution Explain Our Reliability about Logic? *Oxford Studies in Epistemology*, 4:214–23.

Warren, J. (2015). Talking with Tonkers. *Philosophers' Imprint*, 15(24).

Williamson, T. (forthcoming). Modal Science. *The Philosophy of Modality: Beyond the Standard Model*.

Woods, J. (2012). Failures of Categoricity and Compositionality for Intuitionistic Disjunction. *Thought: A Journal of Philosophy*, 1(4):281–291.

Woods, J. (forthcoming). Assertion, Denial, Content, and (logical) Form. *Synthese*.

Woods, J. (Manuscript). Logical Neutrality: A Plea for Ecumenical Reasons.