Penultimate draft. Please cite the final version published in *Topoi*.

Implication Space Semantics as Bilateral Incompatibility Semantics

Ryan Simonelli

April 10, 2026

Abstract

Robert Brandom has described "implication space" semantics, first put forward in Daniel Kaplan's dissertation, as the "holy grail" of inferentialist semantics. However, when one tries to work through the details of the semantic framework, it is very hard to make intuitive sense of what the semantic clauses actually say. In this paper, I articulate a novel interpretation of implication space semantics as a bilateral successor to Brandom's earlier "incompatibility semantics." On this interpretation of the framework, semantic values are assigned, in the first instance, to "positions" consisting in assertions and denials, and the semantic value of such a position is the set of positions that are incompatible with it. I show how this interpretation renders otherwise obscure features of this semantic framework intuitively intelligible in terms of notions previously articulated by Brandom, most notably, that of incompatibility entailment.

Keywords: semantic inferentialism; logical expressivism; bilateralism; incompatibility semantics; implication space semantics

1 Introduction

"Implication space semantics," first put forward in Daniel Kaplan's [13] dissertation, is the first working compositional inferentialist semantic framework, specifying the inferential roles of sentences in terms of the inferential roles of their components, and doing so in such a way as to permit radically substructural inferential relations, the inclusion of which is essential to an empirically adequate inferentialist theory. Robert Brandom [2] describes this framework as what "inferentialists have always dreamed about." However, though the basic idea of the framework is straightforward enough, when one tries to work through the details, it is nearly impossible to understand what the semantic clauses actually say. In Hlobil and Brandom's [10] recent *Reason for Logic, Logic for Reasons* (henceforth *RLLR*), they present a version of the framework that moves to a higher level of abstraction, only compounding the core problem of the basic intelligibility of the framework. In this paper, I present a novel interpretation of Kaplan's original "implication space semantics" as a bilateral successor to Brandom's [1] "incompatibility semantics." On this interpretation of the framework, inspired by Restall's [20] bilateral reading of consequence to which Hlobil and Brandom appeal, semantic values are assigned, in the first instance, to "positions" consisting in assertions and denials, and the semantic value of such a position is the set of positions that are incompatible with it. I show how this interpretation enables one to make clear sense of the otherwise obscure features of this semantic framework in terms of notions previously articulated by Brandom, most notably, that of incompatibility entailment.

2 The Current State of "The Current State of the Art"

Brandom's landmark *Making It Explicit*, published in 1994, systematically lays out inferentialism as a global theory of meaning. However, it contains no formal framework for actually doing inferentialist semantics.[1] Brandom's first real attempt at such a formal framework didn't come until the formal incompatibility semantics put forward in his 2006 Locke Lectures (published in 2008 as *Between Saying and Doing*). In this framework, the meaning of a sentence is understood in terms of the sets of sentences with which that sentence is incompatible. This framework, however, had a crucial problem: it only worked on the assumption that incompatibility relations were *persistent*. That is, it was built into the semantic framework at ground level that if a set of sentences X is incompatible with a set of sentences Y, then any superset of X is incompatible with Y. The problem is that the concept of material incompatibility that the

[1] One step towards a formal inferentialist framework was made by Lance and Kremer [16] (see also [17] and [15]), who, the same year as *Making It Explicit* was published, put forward a set of proof systems for conditionals meant to capture the notion of committive consequence that plays a central role in Brandom's work. These systems were quite limited in scope, however, with really just the conditional as the target connective, and didn't offer the prospect of a general framework for inferentialist semantics.

semantics is meant to be modeling simply doesn't work like this. For instance, "Sadie's a mammal" is incompatible with "Sadie lays eggs," but "Sadie's a mammal" along with "Sadie's a platypus" isn't incompatible with "Sadie lays eggs." Given the prevalence of such defeasible incompatibility relations in natural language, this is a serious problem.[2]

It is for this reason that, in the early 2010s, Brandom, moved by technical work by his student Ulf Hlobil, ended up coming around to formulating inferentialism in terms of the sequent calculus.[3] One feature of the sequent calculus is that it forces one to explicitly use structural rules such as Monotonicity, or, as Gentzen put it, *Weakening*:

$$\frac{X \vdash A}{X, B \vdash A}\ \text{Weakening}$$

The fact that the use of a structural rule such as Weakening itself constitutes a logical step in the sequent calculus makes it possible to construct *substructural logics*: logical systems that work without the use of such rules. Now, there are different reasons to want a logical system that works without such rules, but, for Brandom, the reason is so that the system is able to accommodate sequents for which they actually fail. Of course, Weakening holds for any strictly *logical* inference. If A logically entails B, then, no matter what premises you add to A, you'll still have a logical entailment. However, by rejecting Weakening, we can introduce into our logical system not just sequents encoding logical entailments, but sequents encoding defeasible material inferential relations as well. For instance, we can add, as a non-logical axiom of our sequent calculus, a sequent such as:

$$\mathbf{bird} \vdash \mathbf{flies}$$

and we can do this while maintaining

$$\mathbf{bird}, \mathbf{penguin} \nvdash \mathbf{flies}$$

[2] See Nickel [18] for a criticism of this framework on these grounds.
[3] For Hlobil's statements of this approach to inferentialism, see [6], [7], and [8].

In this way, a proof-theoretic approach to inferentialism based on the sequent calculus is much better suited to accommodate defeasible inferential relations than the sort of natural deduction approach that is common in proof-theoretic semantics.

Now Gentzen's own sequent calculi require the structural rule of Weakening to function. Moreover, the connective rules enforce Weakening with conjuncts and disjuncts. For instance, Gentzen's left-conjunction rules are the following:

$$\frac{X, A \vdash Y}{X, A \wedge B \vdash Y}\ \mathrm{L}_{\wedge_1} \qquad\qquad \frac{X, B \vdash Y}{X, A \wedge B \vdash Y}\ \mathrm{L}_{\wedge_2}$$

These rules would let us reason from

bird $\vdash$ **flies**

to

bird $\wedge$ **penguin** $\vdash$ **flies**.

And, of course, this is an unacceptable consequence. However, by tweaking the rules, we can avoid such consequences. Kaplan's [12] first major contribution to the formal development of inferentialism was showing that by using the classical sequent calculus put forward by Oiva Ketonen [14] in his 1944 dissertation, these unacceptable consequences are elegantly avoided. Here is Ketonen's classical sequent calculus:

$$\frac{}{X, A \vdash A, Y}\ \mathrm{Ax}$$

Where X, Y, and $\{A\}$ contain only atomics.

$$\frac{X \vdash A, Y}{X, \neg A \vdash Y}\ \mathrm{L}_{\neg} \qquad\qquad \frac{X, A \vdash Y}{X \vdash \neg A, Y}\ \mathrm{R}_{\neg}$$

$$\frac{X, A, B \vdash Y}{X, A \wedge B \vdash Y}\ \mathrm{L}_{\wedge} \qquad\qquad \frac{X \vdash A, Y \quad X \vdash B, Y}{X \vdash A \wedge B, Y}\ \mathrm{R}_{\wedge}$$

$$\frac{X, A \vdash Y \quad X, B \vdash Y}{X, A \vee B \vdash Y}\ \mathrm{L}_{\vee} \qquad\qquad \frac{X \vdash A, B, Y}{X \vdash A \vee B, Y}\ \mathrm{R}_{\vee}$$

$$\frac{X \vdash A, Y \quad X, B \vdash Y}{X, A \to B \vdash Y}\ \mathrm{L}_{\to} \qquad\qquad \frac{X, A \vdash B, Y}{X \vdash A \to B, Y}\ \mathrm{R}_{\to}$$

To arrive at a formal proof-theoretic approach approach to inferentialism, one can simply add to this sequent calculus additional non-logical material axioms for which Weakening may actually fail, such as the following:

1. **red** $\vdash$ **colored**
2. **red**, **green** $\vdash$
3. **bird** $\vdash$ **flies**
4. **mammal**, **lays eggs** $\vdash$

Weakening holds for the first two of these sequents, but it fails for the second two, and, notably, all of them are integrated into the same logical system whose rules are proposed as definitive of the meanings of the logical connectives.

Doing things in terms of the sequent calculus in this way constituted a definitive advance in the formal development of inferentialism over Brandom's earlier way of doing things in that the approach now enables the accommodation of defeasible material inferential relations. The above development, however, is restricted to understanding the meanings of sentences *proof-theoretically*. While some inferentialists have thought that a proof-theoretic specification of meanings is just what inferentialists ought to content themselves with, one might hold out hope for a *model-theoretic* inferentialist specification of meanings, one that rivals standard representationalist model-theoretic approaches. Indeed, there is reason to be optimistic here. The sequent calculus enables us to start with a *base* consequence relation—a set of sequents encoding basic material inferences featuring only atomic sentences, determining their semantic significance—and use the left and right rules to recursively generate an *extended* consequence relation, determining the semantic significance of all logically complex sentences. Insofar as we think of the semantic value of a sentence in terms of the sets of provable sequents in which it figures, this means that the semantic values of complex sentences belonging to the language are completely determined by the semantic values of simpler sentences. So, it should be possible to provide semantic clauses that specify, directly, how the semantic values of complex sentences are determined by the semantic values of simpler ones. This is what the "implication space semantics," put forward by Kaplan in his 2022 dissertation, managed to do.

By Brandom's lights, Kaplan's implication space semantics was a watershed moment in the development of inferentialism. Here is how Brandom introduces this formal inferentialist framework:

> The basic idea of semantic inferentialism is to understand conceptual content in terms of role in implications and incompatibilities. So, right from the beginning, the most sought after prize of the inferentialist program—its grail, the one far-off divine event towards which the whole creation moves—has been giving a *direct* specification of the claimables that are expressed by declarative sentences in terms of the relations of implication incompatibility that they stand in to one another. And the criteria of adequacy for that are that a formal inferentialist semantics, in terms of reason relations, has to be as flexible, expressively powerful, and mathematically tractable as the best representational model-theoretic specifications of content. And if I can wax autobiographical for a minute, I can say, I spent my entire professional career since already in my dissertation looking for something like this, trying to make something like this work. And I feel like I learned a lot along the way, but I never figured out how to do it. *Dan did.* [. . .] His implication space semantics is what we inferentialists have always dreamed about [2, 17:28].

The basic idea of Kaplan's semantics is that, what is to be interpreted, in the first instance, are *candidate implications*. That is, the basic objects to which the semantics assigns values are *candidate* implications of the form $X \vdash Y$, and the semantic value assigned to such a thing is the set of other candidate implications $X' \vdash Y'$, such that $X \cup X' \vdash Y \cup Y'$ is a *good* implication. For an implication that is already good, these sets of candidate implications can be understood as their *ranges of subjunctive robustness* (cf. [1, 104-105]), but this notion is generalized to implications that aren't already good, being the sets of premises and conclusions that would *make them* good. Of particular note are the candidate implications of the form $p \vdash$ and $\vdash p$. Kaplan's framework enables us to assign these candidate implications, respectively, the set of candidate implications $X \vdash Y$ such that $X, p \vdash Y$ is a good implication and the set of candidate implications $X \vdash Y$ such that $X \vdash p, Y$ is a good implication. These two sets can be understood as respectively codifying p's role as a premise and p's role as a conclusion, and the pair of them serves as the semantic value of an atomic sentence p. Kaplan then provides semantic clauses that enable us to recursively specify the semantic values of logically complex sentences, in terms of their role as premise and conclusion, so understood. This semantics, Brandom claims, is what "inferentialists have always dreamed about." However, when one actually dives into the semantics and tries to understand what's actually going on, the dream quickly becomes a nightmare.

While the basic idea of the implication space semantics is straightforward enough, and it's not too hard to see how, mathematically, it's sound and complete with respect to Ketonen's sequent calculus, when it comes to actually trying to understand the formal semantics as providing a specification of the meaning of, say, a simple conjunctive sentence in inferentialist terms, one quickly gets the experience of losing one's grip on one's intuitive understanding of what the mathematical entities that one's using to model this meaning actually are. Indeed, despite Brandom's above expressed enthusiasm for the semantics, the ensuing remarks he makes about the details suggest that this is his experience as well. There are, I think, two intersecting issues that make the semantics nearly impossible to comprehend. The first issue is that multiple conclusion implications (which may or may not be *good* implications) serve as the "points" in the semantics, analogous to how possible worlds serve as points in a standard representationalist semantics. Unlike possible worlds, however, (and unlike single conclusion good implications), multiple conclusion implications, which may or may not be good, are not particularly intuitive objects. In the semantics, we are primarily dealing with sets of such candidate implications, operating on them to yield other sets of candidate implications, and, in the context of such operations, it is easy to lose an intuitive grip of what we're actually dealing with. The second compounding issue has to do with the operation of taking a candidate implication's range of subjunctive robustness. As described above, this basic operation is fairly intuitive. However, Kaplan generalizes this operation to apply to *sets* of candidate implications, and he makes essential use of the idea of applying this operation to a set of candidate implications *successively*, taking the set's "range of subjunctive robustness," and then taking the "range of subjunctive robustness" of that range of subjunctive robustness. Such successive applications of this operation play an essential technical role in the framework, but, it's not at all clear what intuitively corresponds to sets we actually end up with when we apply this operation to some set of candidate implications multiple times. The result is a set of semantic clauses and a definition of entailment which, though provably sound and complete relative to the multiple conclusion sequent calculus, are, to put it mildly, far from conceptually transparent.

Now, the most recent chapter in this story is Hlobil and Brandom's recent publication of *RLLR*. Though Kaplan himself is not an author of the book (as originally intended), it is Kaplan's implication space semantics that, as Brandom says, represents "the current state of the art in inferentialist semantics" [10, 18], and it is his semantic framework to which the book is advertised as leading as its culminating moment. In the introduction, Brandom describes the semantics as I've just described it above, saying that, on an implication space semantics, "The range of subjunctive robustness of a candidate implication is [a candidate implication's] semantic interpretant." Ranges of subjective robustness are, indeed, the semantic interpretants of candidate implications in Kaplan's semantics. However, when one gets to Chapter 5 of the book, where the "implication space semantics" is actually presented by Hlobil, the semantic values we get are not Kaplan's semantic values. Rather than assigning to a candidate implication its range of subjunctive robustness as its semantic values, Hlobil assigns to a candidate implication the equivalence class of sets of candidate implications that have the same subjective robustness (in the extended sense in which we can speak of "ranges of subjunctive robustness" of *sets* of candidate implications). Thus, the semantic values that are assigned in Hlobil's version of the framework are not sets of candidate implications but sets of sets of candidate implications. Though, moving to this further level of abstraction, Hlobil states semantic clauses that are visually simpler, comprehending what actually corresponds to the semantic values yielded by these clauses is even more difficult. Even bracketing a problem with these clauses that I outline (and resolve) in the appendix below, getting a concrete grip on what the semantic value of even a simple sentence such as "*a* is red and round" is actually supposed to be, in Hlobil's version of implication space semantics, can seem outright impossible.[4]

All of this leads one to wonder: if this is indeed the "grail" that whole inferentialist program has been forever searching for, then perhaps so much the worse for the whole program. Less radically, perhaps the aspiration for a compositional model-theoretic inferentialist semantics is misguided, and inferentialists should simply content themselves with recursive proof-theoretic specifications of inferential roles. This latter view is, in fact, the view towards which I mostly find myself inclined. Nevertheless, I still think that Kaplan's implication space semantics deserves a better presentation than the one it has been given, either in Kaplan's own work or in Hlobil and Brandom's. That's what I'll aim to do here.

3 Going Explicitly Bilateral

I will start by resolving the first major issue stated above, involving the basic points in the semantics being candidate multiple conclusion implications. It's perhaps worth being explicit about why this is a major issue. The crucial idea of a multiple conclusion implication is that the premises, collectively, "imply" the conclusions, collectively; however, whereas the premises are collected *conjunctively*, the conclusions are collected *disjunctively*. Such multiple conclusion "implications" are theoretical objects that do not find much pre-theoretical

[4]An anonymous referee wonders what weight such epistemic considerations should actually bear on the question of the correct semantic theory. Might it just be the case that the correct semantic theory is very difficult to understand, just as, perhaps, the correct theory of quantum mechanics is very difficult to understand? Indeed, might it be that, just as the correct theory of quantum mechanics might be impossible to understand by creatures with the limited cognitive capacities we happen to have, so too is the correct semantic theory? I think not, as I take it that there is an important disanalogy between the two disciplines. Whereas the correct theory of quantum physics has essentially no relation to our own cognitive capacities, and thus, theorizing about quantum physics is not bound by a strong "intelligibility constraint," a semantic theory aims to *explicitly explicate* what speakers *already implicitly grasp* in knowing how to use the expressions of a language they've mastered. Thus, the theory should at least be such that it can be made intelligible to those very speakers who already implicitly grasp it.

traction in our actual practices of reasoning. Of course, we have a grip of what it is for a set of premises to imply a single conclusion that is a *disjunction*, but the multiple conclusions of a multiple conclusion implication are not to be interpreted in that way any more than the multiple premises of a single conclusion implication are to be interpreted as a single conjunction, as doing so would preclude us from being able to appeal to such implications in giving an account of the meanings of these propositional connectives.

The issue of making good sense of multiple conclusion implications—and doing so in a way that does not presuppose grasp of the logical connectives whose meanings they are supposed to be formally accounting for—has long been the bugbear haunting proponents of multiple conclusion frameworks in the context of inferentialist semantics. Many proponents of inferentialist semantics remain convinced that, as Florian Steinberger [28] puts it, "Conclusions Should Remain Single," leading many to prefer single conclusion systems (typically, natural deduction systems) in formally developing an inferentialist approach to meaning. Indeed, Brandom himself has been no exception here, and his developments of inferentialism prior to Kaplan's technical developments clearly show a preference for the single conclusion setting. In personal correspondence, he has described Kaplan, who's technical developments centrally rely on sequents having multiple conclusion, as having "dragged him kicking and screaming" towards a multiple conclusion setting. Eventually, the need to make sense of multiple conclusion sequents led Brandom (after some initial resistance) to settle on a philosophical approach to multiple conclusion sequents owed to Greg Restall and Ellie Ripley

Restall [20] proposes a reading of multiple conclusion sequents according to which the turnstile plays the role not of separating *premises* from *conclusions* but of separating *assertions* from *denials*. It is this *bilateral* approach, developed by Restall, that Hlobil and Brandom [10, 106] officially adopt in presenting their formal inferentialist theory of meaning, couched in terms of multiple conclusion sequents. On this bilateral reading of multiple conclusion sequents, a multiple conclusion sequent of the form $X \vdash Y$ is read as saying that the position, $X : Y$, consisting in asserting everything in X and denying everything in Y is incoherent, "out of bounds" [21], or involves some sort of "clash."[5] To illustrate how this reading resolves these concerns, consider again the negation rules of the multiple conclusion sequent calculus:

$$\frac{X \vdash A, Y}{X, \neg A \vdash Y}\ \mathrm{L}_{\neg} \qquad\qquad \frac{X, A \vdash Y}{X \vdash \neg A, Y}\ \mathrm{R}_{\neg}$$

While the meaning of these rules will seem opaque if the horizontal lines are read in the usual way as deductively relating *implications*, Restall's bilateral reading renders their meaning perfectly transparent by reading the horizontal lines as relating *incoherences.* On Restall's proposed reading, the left rule simply says that if, relative to any position $X : Y$, denying A is incoherent, then, relative to $X : Y$, asserting $\neg A$ is incoherent. Similarly, the right rule says that if, relative to any position $X : Y$, asserting A is incoherent, then, relative to $X : Y$, denying $\neg A$ is incoherent. Thus, understanding the significance of assertions and denials in terms of their potential contribution to the incoherence of a discursive position, these rules together say that asserting the negation of some sentence has the same significance as denying that sentence, and denying the negation of some sentence has the same significance as asserting that sentence.

If we think that specifying a sentence's role in a sequent calculus, so understood, suffices to specify its meaning, then the general principle for understanding the meaning of a sentence, on a Restall-style bilateralism, might be put as follows:

> *The meaning of a sentence is the contribution that its assertion and denial makes to the incoherence of positions.*

I suggest that we take this principle seriously, and think of an inferentialist theory articulated in terms of multiple conclusion sequents as a kind of *incompatibility* semantics, of the sort proposed by Brandom [1]. However, rather than a *uni*lateral incompatibility semantics, where we understand the semantic significance of a sentence in terms of the sets of other sentences with which

[5] Thus, for instance, it is surely not, in the strict sense, *incoherent* to assert that something's a bird and (without giving any further information) deny that it flies. However, this combination of moves might nevertheless be understood as involving a sort of "tension" or "clash," calling out for more moves to be made (such as an assertion of "It's a penguin") to resolve it.

it is incompatible, we'll have a *bi*lateral incompatibility semantics, where we understand the semantic significance of an *assertion* or a *denial* of a sentence in terms of the sets of other *assertions and denials* with which it is incompatible. I suggest that this is how Kaplan's implication space semantics is best understood. Thus, rather than the "points" of the semantics being candidate multiple conclusion implications, we understand them simply as possible positions that one might occupy, where a position is just any set of assertions and denials. Let us now spell this idea out officially.

4 Incompatibility Semantics, Redux

We start with a language, $\mathcal{L}$, and define the following semantic ingredients:

Definition 1 (Moves): The total set of *moves*, $\mathcal{L}_{\pm}$, is $\{+A \mid A \in \mathcal{L}\} \cup \{-A \mid A \in \mathcal{L}\}$.

Definition 2 (Positions): The total set of *positions*, $\mathbb{P}$, is the powerset of $\mathcal{L}_{\pm}$.

Definition 2.1 The *minimal* position, e, is $\varnothing$
Definition 2.2 The *maximal* position, $\star$, is $\mathcal{L}_{\pm}$.

Definition 3 (Incoherent Positions): There is a distinguished subset of *incoherent* positions, $\mathbb{I} \subseteq \mathbb{P}$.[6]

Constraint 1: $e \notin \mathbb{I}$ and $\star \in \mathbb{I}$

[6] One might be interested in imposing further structure on the set of incoherences. For instance, one plausible further constraint is the following:

Extensibility: For any position Γ and any atomic sentence p, if $\Gamma \notin \mathbb{I}$, then $\Gamma \cup \{+p\} \notin \mathbb{I}$ or $\Gamma \cup \{-p\} \notin \mathbb{I}$

This principle corresponds to the shared context version of *Cut* in the multiple conclusion sequent calculus. Now, in the context of an inferentialist theory that accommodates defeasible inferences, there are obvious issues with the mixed context version of Cut, and Hlobil [6] and myself [22] [27] present arguments against the shared context version of Cut in a *single conclusion* context. However, the shared context version of Cut in a *multiple* conclusion context—discussed by Restall [20] and Ripley [21] as a principle of "extensibility" (and which Ripley rejects for reasons having to do with semantic paradoxes)—may plausibly be accepted in this context. Following Kaplan, Hlobil, and Brandom's presentation, I will not impose this additional structure, but it is worth noting that it could in principle be imposed.

Constraint 2: For any position Γ, and any atomic sentence p, $\Gamma \cup \{+p, -p\} \in \mathbb{I}$.

So, for a given language, the assertion of any sentence of the language is a move, and the denial of any sentence of the language is a move. A position is any set of moves. The minimal position consists in neither asserting nor denying anything, and the maximal position consists in both asserting and denying everything. Among the total set of positions is a distinguished subset of positions that are incoherent. We assume that the minimal position is coherent, the maximal position is incoherent, and, moreover, that any position that contains the assertion and the denial of the same atomic sentence is incoherent. In addition to such *formal* incoherences, imposed by the framework itself, we will also admit *material* incoherences into $\mathbb{I}$. For instance, if we're giving an incompatibility semantics for English, since asserting "*a* is red" and asserting "*a* is green" is incoherent, we'll have $\{+\mathbf{red}(a), +\mathbf{green}(a)\} \in \mathbb{I}$. Likewise, asserting "*a* is red" and denying "*a* is colored" is incoherent, and so $\{+\mathbf{red}(a), -\mathbf{colored}(a)\} \in \mathbb{I}$. Crucially, we allow that the following principle, central to Brandom's [1] incompatibility semantics, *may fail*:

Persistence: If $\Gamma \subseteq \Delta$, then, if $\Gamma \in \mathbb{I}$, then $\Delta \in \mathbb{I}$.

Thus, we may have, for instance, $\{+\mathbf{mammal}(s), +\mathbf{lays\ eggs}(s)\} \in \mathbb{I}$, but $\{+\mathbf{mammal}(s), +\mathbf{platypus}(s), +\mathbf{lays\ eggs}(s)\} \notin \mathbb{I}$. Likewise, we may have $\{+\mathbf{bird}(b), -\mathbf{flies}(b)\} \in \mathbb{I}$, but $\{+\mathbf{bird}(b), -\mathbf{flies}(b), +\mathbf{penguin}(b)\} \notin \mathbb{I}$

The semantic significance of the assertion or denial of some sentence can be understood in terms of its *incoherence profile*. We might picture the incoherence profile of the assertion of "*a* is red" in the following way:

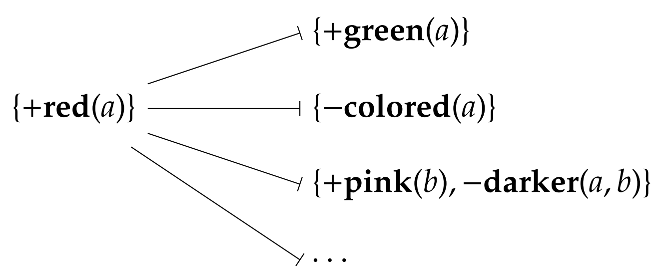

So, on the left, we have the position consisting in asserting "*a* is red," and, on the right, we have all the positions incompatible with asserting "*a* is red": all those positions such that occupying them along with asserting "*a* is red" constitutes an incoherent position. Likewise, we can consider the incoherence profile of the denial of "*a* is red"

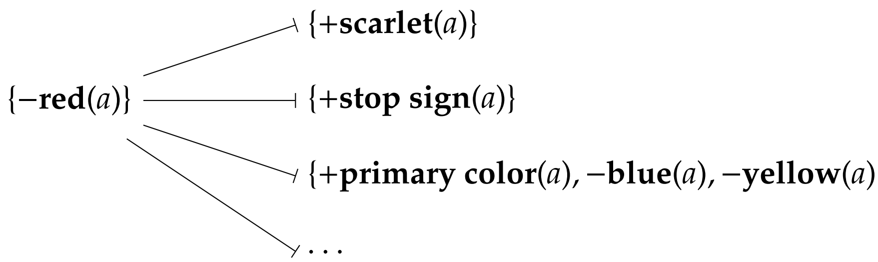

To state these profiles officially, let us define a function $\perp$ which takes some position Γ and returns the set of positions Δ that are incompatible with Γ in the sense that combining Γ with Δ yields an incoherent position:

Definition 4 (Incoherence Profiles): $\Gamma^{\perp} = \{\Delta \mid \Gamma \cup \Delta \in \mathbb{I}\}$

With this function defined, we can understand the significance of asserting "*a* is red" in terms of its incoherence profile, given by $\{+\mathbf{red}(a)\}^{\perp}$.

Though I've just defined incoherence profiles for individual positions, the framework will gain much more expressive power if we generalize the definition of incoherence profiles such that it applies not just to individual positions but also to sets of positions. The incoherence profile of a set of positions is just the intersection of the incoherence profiles of its members. Officially, where **X** is a set of positions, we can define its incoherence profile, $\mathbf{X}^{\perp}$, as follows:

$$\mathbf{X}^{\perp} = \{\Gamma \mid \forall \Delta \in \mathbf{X} : \Gamma \cup \Delta \in \mathbb{I}\}$$

Thus $\mathbf{X}^{\perp}$ is the set of positions Γ such that, for every position $\Delta \in \mathbf{X}$, $\Gamma \cup \Delta$ is an incoherent position. Our previously defined case of taking the incoherence set of a single position, then, is just the special case where **X** is a singleton.

To see the additional expressive power gained by this generalization, consider again our picture of the incoherence set of the position consisting in asserting "*a* is red":

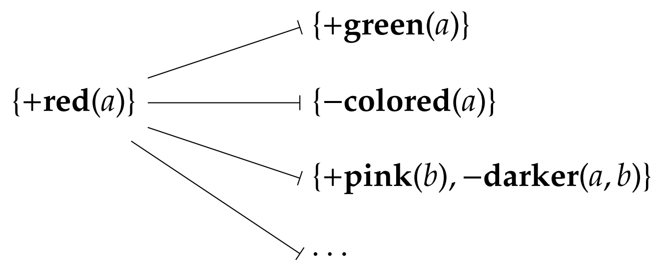

Once again, on the right, we have all positions Δ such that occupying any one of these positions along with asserting "*a* is red" constitutes an incoherent position. A question we may ask about this set on the right is the following: what positions Γ are such that, for any one of these positions on the right (Δ_1, $\Delta_2 \ldots \Delta_n$), $\Gamma \cup \Delta_i$ is incoherent? That is, which positions are in $\{+\mathbf{red}(a)\}^{\perp\perp}$? Clearly, *one* such position is that consisting in asserting "*a* is red." However, it doesn't seem like it's the only possible such position. In particular, any position that has precisely the same inferential significance as asserting "*a* is red" will be such that unioning it with any Δ_i will yield an incoherent position. Moreover, any position that is *inferentially stronger* than asserting "*a* is red," such as asserting "*a* is scarlet," or asserting "*a* is a stop sign," or asserting "*a* is red" along with asserting "*a* is spherical" will be such that unioning it with any Δ_i will yield an incoherent position. Thus, we might picture$\{+\mathbf{red}(a)\}$, $\{+\mathbf{red}(a)\}^{\perp}$, and $\{+\mathbf{red}(a)\}^{\perp\perp}$ as follows:

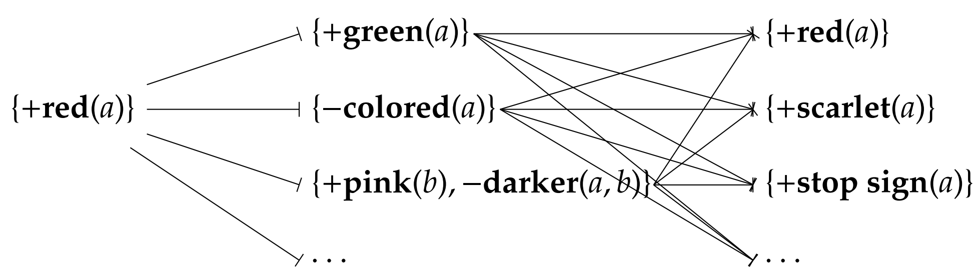

These positions in $\{\mathbf{red}(a)\}^{\perp\perp}$ are such that everything that is incompatible with asserting "*a* is red" is incompatible with each of them. In the terminology of Brandom [1], they each *incompatibility entail* the assertion of "*a* is red."

Given the above picture, one might wonder whether $\{+\mathbf{red}(a)\}^{\perp\perp}$, the set of positions that incompatibility entail the assertion of "*a* is red," is in fact the same as $\{-\mathbf{red}(a)\}^{\perp}$, the set of positions incompatible with denying "*a* is red," and thus, given that we have the $^{\perp}$ function, whether treating assertions and denials independently is really necessary.[7] To see that it is indeed necessary, consider again the example of asserting "Sadie's a mammal." This is incompatible with asserting "Sadie lays eggs." And so $\{+\mathbf{lays\ eggs}(s)\}$ is in $\{+\mathbf{mammal}(s)\}^{\perp}$. Asserting "Sadie's a platypus" is, of course, not incompatible with asserting "Sadie lays eggs." So, $\{+\mathbf{platypus}(s)\}$ is *not* in $\{+\mathbf{mammal}(s)\}^{\perp\perp}$. However, $\{+\mathbf{platypus}(s)\}$ is in $\{-\mathbf{mammal}(s)\}^{\perp}$, since asserting that something's a platypus and denying that it's a mammal is surely incoherent. Thus, $\{+\mathbf{mammal}(s)\}^{\perp\perp} \neq \{-\mathbf{mammal}(s)\}^{\perp}$. The incoherence set of the incoherence set of asserting "Sadie's a mammal" is not the same as the incoherence set of denying "Sadie's a mammal." This shows that, insofar as we want to accommodate these failures of persistence that arise in the context of material inferential relations, we must treat assertion and denial separately.

5 Features of Incompatibility Entailment

It is worth pausing to say a few words on the notion of incompatibility entailment defined in this framework. Officially, let us define this notion of incompatibility entailment as follows:

> **Incompatibility Entailment:** Δ *incompatibility entails* Γ, $\Delta \vDash_I \Gamma$, just in case $\Gamma^{\perp} \subseteq \Delta^{\perp}$

To show that this is indeed the notion we've just defined, we can note the following:[8]

[7]Technically, treating assertion and denial as distinct (making essential use of the two-sidedness of multiple conclusion sequents) is the core difference between Kaplan's implication space semantics and the phase space semantics for linear logic developed by Girard [5] from which it is inspired, allowing implication space semantics to handle radically substructural consequence relations that Girard's phase space semantics cannot. With the use of the $^{\perp}$ function, Girard allows himself to collapse the two sides of the sequent calculus, thus simplifying things by semantically treating only one side of it. In the context of this interpretation, this additional dimension of distinction corresponds to the distinction between a *bilateral* incompatibility semantics and a *unilateral* incompatibility semantics of the sort originally proposed by Brandom. See [19] for a version of unilateral incompatibility semantics based on Girard's phase space semantics for linear logic.

[8]*Proof:* For the left to right direction, note that $\Gamma^{\perp\perp} = \{\Delta \mid \forall\Theta \in \Gamma^{\perp} : \Delta \cup \Theta \in \mathbb{I}\}$. Given that

Fact 1: $\Delta \in \Gamma^{\perp\perp}$ just in case $\Gamma^{\perp} \subseteq \Delta^{\perp}$

Thus, defining incompatibility entailment as above, Δ incompatibility entails Γ just in case everything incompatible with Δ is incompatible with Γ.[9] Likewise, if Δ incompatibility entails Γ and Γ incompatibility entails Δ, then Δ and Γ are *incompatibility equivalent* in that $\Delta^{\perp} = \Gamma^{\perp}$.

Incompatibility entailment is the central semantic notion in Brandom's [1] incompatibility semantics.[10] There are, however, two crucial differences between Brandom's notion of incompatibility entailment and the one defined here. The first is that this is a *bilateral* entailment relation, obtaining not between *sentences*, but between *positions* consisting in assertions and denials. Considering just the case where Γ and Δ are singletons, a move φ incompatibility entails ψ just in case everything incompatible with ψ is incompatible with φ. For instance, asserting "*a* is crimson" incompatibility entails asserting "*a* is red," since everything incompatible with asserting "*a* is red" (e.g. asserting "*a* is green," denying "*a* is colored," and so on) is incompatible with asserting "*a* is crimson." Likewise, asserting "*a* is red" incompatibility entails denying "*a* is green," since everything incompatible with denying "*a* is green" (e.g. asserting "*a* is lime green," asserting "*a* is grass," and so on) is incompatible with asserting "*a* is red." A second and even more significant difference between the framework here and the one endorsed by Brandom there, is that, in this context, we permit *defeasible* incompatibility relations. That is, we permit

$\Gamma^{\perp} = \{\Theta \mid \Gamma \cup \Theta \in \mathbb{I}\}$, it follows that for any $\Delta \in \Gamma^{\perp\perp}$, if $\Gamma \cup \Theta \in \mathbb{I}$, then $\Delta \cup \Theta \in \mathbb{I}$. So, if $\Delta \in \Gamma^{\perp\perp}$, then $\Gamma^{\perp} \subseteq \Delta^{\perp}$. Now consider the right to left direction. We will argue contrapositively. Suppose $\Delta \notin \Gamma^{\perp\perp}$. Then there's some $\Theta \in \Gamma^{\perp}$ such that $\Delta \cup \Theta \notin \mathbb{I}$. It follows directly that $\Gamma^{\perp} \nsubseteq \Delta^{\perp}$. □

[9] Note that the elements of both Γ and Δ are collected conjunctively, so this is not a standard notion of multiple conclusion consequence. Insofar as one wants to explore the consequence relation of incompatibility entailment defined here, it might be helpful to see [4] for an investigation of consequence relations where conclusions are collected conjunctively. It is obviously possible to define a standard multiple conclusion consequence relation for incompatibility entailment in this framework, but I will not pursue that here.

[10] Here is how Brandom explicates the notion:

> *p* incompatibility-entails *q* just in case everything incompatible with *q* is incompatible with *p*. Thus "Pedro is a donkey," incompatibility-entails "Pedro is a mammal," for everything incompatible with Pedro's being a mammal (for instance, Pedro's being an invertebrate, an electronic apparatus, a prime number) is incompatible with Pedro's being a donkey [1, 121].

failure of persistence, defined above. It is straightforward to show that if we impose persistence, consider the fragment of $\mathbb{P}$ consisting solely in sets of positively signed sentences, and restrict consequents to singletons, it aligns with Brandom's incompatibility semantics in this framework. However, because of persistence failures, the notion of incompatibility entailment defined here behaves quite differently.

First, because of persistence failures, the intuitive notion of "implication" applicable in an inferentialist theory (which may be cashed out in terms of (potentially defeasible) committive consequence), on the one hand, and incompatibility entailment, on the other hand, come apart. For instance, asserting "Bella's a bird" commits one to asserting "Bella flies," but asserting "Bella's a bird" does not incompatibility entail asserting "Bella flies," since it's not the case that everything incompatible with asserting "Bella flies" is incompatible with asserting "Bella's a bird." For instance, asserting "Bella's a penguin" is incompatible with asserting "Bella flies," but it's not incompatible with asserting "Bella's a bird." In fact, even in a case of strict implication, there is not necessarily an incompatibility entailment. For example, asserting "Sadie's a platypus" strictly commits one to asserting "Sadie's a mammal," but "Sadie's a platypus" does not incompatibility entail "Sadie's a mammal," since it's not the case that everything incompatible with "Sadie's a mammal" (e.g. "Sadie lays eggs") is incompatible with "Sadie is a platypus."

Second, persistence failures have striking consequences for the structural features of incompatibility entailment. For instance, it will not always be the case that $\{+p, +q\}$ incompatibility entails $\{+p\}$, since it may not be the case that everything incompatible with $\{+p\}$ is incompatible with $\{+p, +q\}$. For instance, asserting "Sadie lays eggs" is incompatible with asserting "Sadie's a mammal," so but asserting "Sadie lays eggs" is not incompatible with asserting "Sadie's a platypus" along with asserting "Sadie's a mammal." Thus, $+\textbf{lays eggs}(s) \in \{+\textbf{mammal}(s)\}^{\perp}$, but $+\textbf{lays eggs}(s) \notin \{+\textbf{mammal}(s), +\textbf{platypus}(s)\}^{\perp}$, and so, $\{+\textbf{mammal}(s), +\textbf{platypus}(s)\} \notin \{+\textbf{mammal}(s)\}^{\perp\perp}$. That is, asserting "Sadie's a mammal" along with asserting "Sadie's a platypus" does not incompatibility entail asserting "Sadie's a mammal." In general, the structural principle of *Containment* does not hold of incompatibility entailment. It's not always the case that $\Gamma, \varphi \vDash_I \varphi$.

These features might seem strange, but they're features of the fact that the set of things that incompatibility entail some position are those whose incoherence profiles are *strictly at least as strong as* that position. This fact will play an important role in the semantic clauses to come. I will leave open further investigation of the logic of this notion of incompatibility entailment for future work, turning now to the main fruits of this semantic framework: semantic values.

6 Semantic Values

Semantic values are formal models of the meanings of sentences. In a standard, representational, model-theoretic semantic framework, semantic values are set-theoretic entities constructed out of a domain of *extra-linguistic* entities. For instance, in a standard extensional semantics of the sort proposed by Heim and Kratzer, the semantic value of "Mars" will be the planet, Mars. Or, in an intensional semantics, the semantic value of "Mars is red" will be the set of possible worlds in which Mars is red. An inferentialist semantic framework, however, constructs semantic values entirely in *intra-linguistic* terms. A semantic value will be a set of positions, where these are sets of signed sentences. Now, there is a range of different possible types of semantic values that might be defined in this framework. I will suppose, in keeping with Brandom's [2] advertisement of the framework as comparable to "the best representational model-theoretic specifications of content," that any definition of semantic values must meet the following requirement.

> **Compositionality Requirement:** The semantic value of a complex sentence must be defined in terms of the semantic values of the simpler sentences that compose it.

Thus, if we take semantic values to be incoherence profiles, the semantic clauses must tell us how to compute the incoherence profile of a conjunction, given the incoherence profiles of the conjuncts. I will talk more about this requirement shortly. First, however, let me say some general things about semantic values, as defined in this framework.

A notable feature of this framework is that not only are semantic values sets of positions, but, moreover, semantic values are assigned, in the first instance, *to* positions, and assigned derivatively to sentences as the pair consisting in the assertion of that sentence along with its denial. The two most plausible candidates for the semantic value of a position are (1) its incoherence profile or (2) the incoherence profile of its incoherence profile (i.e. the set of positions that incompatibility entail it).[11] There are things to be said in favor of both approaches. On the one hand semantic values of the first sort are, in some sense, more conceptually basic, and more naturally fit the idea of this being an *incompatibility* semantics. On the other hand, the actual semantic clauses are a bit cleaner if we define semantic values of the second sort, and the idea of taking the semantic value of (an assertion of) a sentence to be the set of things that incompatibility entail actually gives us something rather like truth-conditions, which may be helpful in connecting inferentialism to more familiar truth-conditional theories down the line.[12] Both approaches, I think, are worth exploring. However, following Kaplan's original approach, here I'll define semantic values of the first sort.

For any sentence A whose semantic value is defined in this framework, its semantic value will be the pair consisting, first, in the semantic value of its assertion, and, second, in the semantic value of its denial. That is:

[11]Note, neither of these candidates are the one that Hlobil goes in for. See the appendix for a critical consideration of Hlobil's alternative approach.

[12]Though I won't justify these semantic clauses here, they are as follows (the $\Cup$ operation is defined below):

$$\mathbf{S}_{\mathcal{A}} : [\![p]\!] = \langle \{+p\}^{\perp\perp}, \{-p\}^{\perp\perp} \rangle$$
$$\mathbf{S}_{\wedge} : [\![A \wedge B]\!] = \langle ([\![+A]\!] \Cup [\![+B]\!])^{\perp\perp}, [\![-A]\!] \cup [\![-B]\!] \rangle$$
$$\mathbf{S}_{\vee} : [\![A \vee B]\!] = \langle [\![+A]\!] \cup [\![+B]\!], ([\![-A]\!] \Cup [\![-B]\!])^{\perp\perp} \rangle$$
$$\mathbf{S}_{\rightarrow} : [\![A \rightarrow B]\!] = \langle [\![-A]\!] \cup [\![+B]\!], ([\![+A]\!] \Cup [\![-B]\!])^{\perp\perp} \rangle$$

The notion of entailment (for just the two sentence case) is as follows:

Entailment: $A \vDash B$ just in case $([\![+A]\!] \Cup [\![-B]\!])^{\perp\perp} \subseteq \mathbb{I}$

In ongoing work developing a category-theoretic formalization of implication space semantics, Kris Brown [3] has been working with semantic values of this sort.

$$[\![A]\!] = \langle [\![+A]\!], [\![-A]\!] \rangle$$

These two semantic values are, respectively, the incoherence profile of the assertion of A and the incoherence profile of the denial of A. For atomic sentences, we take such specifications of profiles as basic, directly given by the function $^{\perp}$. So, for any atomic sentence p, $[\![+p]\!] = \{+p\}^{\perp}$ and $[\![-p]\!] = \{-p\}^{\perp}$. Thus, we have:

$$\mathbf{S}_{\mathcal{A}} : [\![p]\!] = \langle \{+p\}^{\perp}, \{-p\}^{\perp} \rangle$$

For logically complex sentences, we want to specify recursive semantic clauses such that, given the incoherence profiles of the simpler sentences that compose them, we can determine the incoherence profile of the complex sentence.

Negation is straightforward, given the bilateralist understanding of it as an operator that flips between assertion and denial. So, the profile of asserting $\neg A$ is the same as that of denying A, and the profile of denying $\neg A$ is the same as that of asserting A. Thus, the semantic clause for negation is as follows:[13]

$$\mathbf{S}_{\neg}\text{: } [\![\neg A]\!] = \langle [\![-A]\!], [\![+A]\!] \rangle$$

It is easy to see that $[\![A]\!] = [\![\neg\neg A]\!]$.

Let us now consider the profile of asserting or denying a conjunction. Consider first the profile of denying $A \wedge B$. Which positions are such that, occupying them, denying $A \wedge B$ is incoherent? Intuitively, if I can coherently deny A or I can coherently deny B, then I can coherently deny $A \wedge B$. So, the positions relative to which denying $A \wedge B$ is *in*coherent are just those such that, relative to them, denying A is incoherent and denying B is also incoherent. Thus, the incoherence profile of denying $A \wedge B$ is the intersection of that of denying A and that of denying B:

$$[\![-A \wedge B]\!] = [\![-A]\!] \cap [\![-B]\!]$$

[13] This sort of semantic clause for negation will be familiar to those acquainted with truth-maker semantics of the sort proposed by Fine (2017), and the semantic clauses for the binary connectives to follow also bear some similarity. For a discussion of the relation between implication space semantics and Fine-style truth-maker semantics, see Hlobil and Brandom (2024, 221-228).

What about the profile of asserting $A \wedge B$? Which positions are such that, occupying them, asserting $A \wedge B$ is incoherent? Intuitively, these are just the positions such that occupying them, asserting A along with asserting B is incoherent. Thus, we should want the profile of asserting $A \wedge B$ to come out as follows:

$$[\![+A \wedge B]\!] = \{+A, +B\}^{\perp}$$

However, this doesn't suffice as a semantic clause, since it does not satisfy the compositionality requirement, which requires us to define $[\![+A \wedge B]\!]$ in terms of $[\![A]\!]$ and $[\![B]\!]$. Before showing how this is done, let me pause to say a few words about the compositionality requirement to clarify what it actually amounts to.

Brief Interlude: Compositionality vs. Recursivity

In order to clarify the compositionality requirement, let us contrast it with the following, weaker requirement:

> **Recursivity Requirement:** The semantic value of a complex sentence must be determined by the semantic values of the simpler sentences that compose it.

One might think that it is sufficient to specify semantic clauses that meet this requirement. However, if we're going to settle for this requirement, it's hard to see how model-theoretic semantics would really distinguish itself from a proof-theoretic approach provided by the sequent calculus. We already know that the sequent calculus recursively determines the valid sequents in which a logically complex sentence figures, given the valid sequents in which the simpler sentences that compose it figure. Because of this feature of the sequent calculus, the rules of the sequent calculus can *already* be understood as semantic clauses that satisfy the recursivity requirement. Consider the following sequent rules:

$$\frac{X, A, B \vdash Y}{X, A \wedge B \vdash Y}\ L_{\wedge} \qquad\qquad \frac{X \vdash A, Y \quad X \vdash B, Y}{X \vdash A \wedge B, Y}\ R_{\wedge}$$

On the bilateralist reading, the left rule is understood as saying, relative to any context consisting in asserting everything in X along with denying everything in Y, if asserting A along with asserting B is incoherent, then, relative to that context, asserting $A \wedge B$ is incoherent. Notably, these rules are invertible, and so we can say, moreover, that these are *precisely* the contexts relative to which asserting $A \wedge B$ is incoherent. Likewise, the right rule says, relative to any context consisting in asserting everything in X along with denying everything in Y, if denying A is incoherent and, relative to that same context, denying B is incoherent, then, relative to that context, denying $A \wedge B$ is incoherent. Once again, given the invertibility of the rules, we can say that these are precisely the contexts relative to which denying $A \wedge B$ is incoherent. Now, mapping a multiple conclusion sequent of the form $X \vdash Y$ to a position $\Gamma = \{+A \mid A \in X\} \cup \{-B \mid B \in Y\}$, and taking a valid sequent to be one such that $\Gamma \in \mathbb{I}$, these left and right sequent rules, so interpreted, give us the following "semantic clause":

$$[\![A \wedge B]\!] = \langle \{\Gamma \mid \{\Gamma, +A, +B\} \in \mathbb{I}\}, \{\Gamma \mid \{\Gamma, -A\} \in \mathbb{I} \text{ and } \{\Gamma, -B\} \in \mathbb{I}\} \rangle$$

The first element of the pair specifies the contexts that figure in the top sequent of the left conjunction rule, whereas the second element of the pair specifies the contexts that figure in the top two sequents of the right conjunction rule. Given our definition of $^{\perp}$, we can rewrite this as the following:

$$[\![A \wedge B]\!] = \langle \{+A, +B\}^{\perp}, \{-A\}^{\perp} \cap \{-B\}^{\perp} \rangle$$

We can provide similar clauses for all of the other connectives, based on their sequent rules. Giving such clauses—which essentially just amounts to rewriting the rules of the sequent calculus—provides a recursive specification of the meanings of logically complex sentences in terms of the positions that are incompatible with them. However, though the second element of this pair specifies an operation on the *semantic values* of denying A and denying B, the first element of this pair does no such thing. Accordingly, such clauses don't actually define the semantic value of a conjunction in terms of operations on the set-theoretic entities that are the semantic values of the conjuncts. In this way, though these clauses meet the recursivity requirement, they do not meet the compositionality requirement.

For the purposes of this paper, I have bracketed the question of to what extent it is actually *important* for an inferentialist semantics to meet the compositionality requirement. This, in effect, is just the question of whether an inferentialist ought to want a compositional, model-theoretic semantics or can be content with a recursive proof-theoretic semantics. As I have already indicated, my own sympathies primarily lie with recursive proof-theoretic semantics. However, if we do want a compositional, model-theoretic semantics, at least the *criterion* we need to meet is clear. Let us now see how we can meet it.

Semantic Values, Continued

We return now to the case of defining the semantic value of asserting a conjunction. Once again, we need to define an operation on $[\![+A]\!]$ and $[\![+B]\!]$ such that the result of the operation is the same as $\{+A, +B\}^{\perp}$. Here's an idea. Consider $[\![+A]\!]^{\perp}$ and $[\![+B]\!]^{\perp}$. These, once again, are $\{+A\}^{\perp\perp}$ and $\{+B\}^{\perp\perp}$, which are, respectively, the set of positions that incompatibility entail $+A$ and the set of positions that incompatibility entail $+B$. Each position from the first set is at least as strong as $+A$ and each position from the second set is at least as strong as $+B$, and so combining them yields a position at least as strong as $+A$ and $+B$ together. Now consider the set of positions that are each incompatible with each position in that set. Since $\{+A, +B\}$ is the weakest position in this set (every other position incompatibility entails $\{+A, +B\}$ in the sense that, if something's incompatible with $\{+A, +B\}$, then it's incompatible with that position too), these will be just the positions that are incompatible with $\{+A, +B\}$. To state this idea officially, let us first define an operation that combines sets of positions by unioning pairwise, taking the set consisting in all of the unions of an element of one set of positions along with an element of the other. That is:

$$\mathbf{X} \Cup \mathbf{Y} = \{\Gamma \cup \Delta \mid \Gamma \in \mathbf{X}, \Delta \in \mathbf{Y}\}$$

We can now officially state the idea just stated as follows:

$$[\![+A \wedge B]\!] = ([\![+A]\!]^{\perp} \Cup [\![+B]\!]^{\perp})^{\perp}$$

This clause *does* satisfy the compositionality requirement. To show that this is indeed the result we want, we can note the following important fact:[14]

Fact 2: $(\Gamma^{\perp\perp} \Cup \Delta^{\perp\perp})^{\perp} = (\Gamma \cup \Delta)^{\perp}$

Thus, $[\![+A \wedge B]\!]$ just is $\{+A, +B\}^{\perp}$.

We can now state the full semantic clauses for all of the binary connectives. Putting the two clauses for conjunction together, we get:

$$\mathbf{S}_{\wedge} : [\![A \wedge B]\!] = \langle ([\![+A]\!]^{\perp} \Cup [\![+B]\!]^{\perp})^{\perp}, [\![-A]\!] \cap [\![-B]\!] \rangle$$

Dually, we get the following clause for disjunction:

$$\mathbf{S}_{\vee} : [\![A \vee B]\!] = \langle [\![+A]\!] \cap [\![+B]\!], ([\![-A]\!]^{\perp} \Cup [\![-B]\!]^{\perp})^{\perp} \rangle$$

And, similarly, the following clause for the conditional:

$$\mathbf{S}_{\to} : [\![A \to B]\!] = \langle [\![-A]\!] \cap [\![+B]\!], ([\![+A]\!]^{\perp} \Cup [\![-B]\!]^{\perp})^{\perp} \rangle$$

Of course, these semantic clauses are not as simple or straightforward as the semantic clauses in familiar truth-conditional frameworks. It is, for instance, much more simple and straightforward to say that semantic value of $A \wedge B$ is the intersection of the sets of worlds in which A is true and the set of worlds in which B is true. Compare this with the clause for the semantic value of (the assertion of) a conjunction as stated and explicated here, which, officially, says that the semantic value of asserting a conjunction is the incoherence profile of the pairwise union of the incoherence profile of the incoherence profile of asserting A and the incoherence profile of the incoherence profile of asserting

[14] *Proof:* We'll show first that $(\Gamma^{\perp\perp} \Cup \Delta^{\perp\perp})^{\perp} \subseteq (\Gamma \cup \Delta)^{\perp}$. Suppose $\Theta \in (\Gamma^{\perp\perp} \Cup \Delta^{\perp\perp})^{\perp}$. By definition, for all $\Lambda \in \Gamma^{\perp\perp} \Cup \Delta^{\perp\perp}$, we have $\Theta \cup \Lambda \in \mathbb{I}$. Now, $\Gamma \in \Gamma^{\perp\perp}$ and $\Delta \in \Delta^{\perp\perp}$, so $\Gamma \cup \Delta \in \Gamma^{\perp\perp} \Cup \Delta^{\perp\perp}$. Hence $\Theta \cup \Gamma \cup \Delta \in \mathbb{I}$, which is to say, $\Theta \in (\Gamma \cup \Delta)^{\perp}$. We'll now show that $(\Gamma \cup \Delta)^{\perp} \subseteq (\Gamma^{\perp\perp} \Cup \Delta^{\perp\perp})^{\perp}$. Suppose $\Theta \in (\Gamma \cup \Delta)^{\perp}$. Now consider an arbitrary $\Lambda \in \Gamma^{\perp\perp}$ and $\Sigma \in \Delta^{\perp\perp}$. Since $\Theta \cup \Gamma \cup \Delta \in \mathbb{I}$ and $\Lambda \in \Gamma^{\perp\perp}$, it follows that $\Theta \cup \Lambda \cup \Delta \in \mathbb{I}$. Since $\Sigma \in \Delta^{\perp\perp}$ it follows that $\Theta \cup \Lambda \cup \Sigma \in \mathbb{I}$. Since Λ and Σ are arbitrary, it follows that, for all $\Lambda \in \Gamma^{\perp\perp}$ and $\Sigma \in \Delta^{\perp\perp}$, $\Theta \in (\Lambda \cup \Sigma)^{\perp}$, which is just to say $\Theta \in (\Gamma^{\perp\perp} \Cup \Delta^{\perp\perp})^{\perp}$. □

B. That, of course, is a quite a mouthful, to put it mildly! Still, having interpreted the crucial $^{\perp\perp}$ operation in terms of incompatibility entailment, we at least have a clear account of what this semantic clause actually says, and even just that is a major advancement on prior presentations of this semantic framework. Given the interpretation advanced here, we can gloss this semantic clause as saying that the semantic value of asserting $A \wedge B$ is the incoherence profile of the pairwise union of the set of positions that incompatibility entail A and the set of positions that incompatibility entail B. That, of course, is still not particularly simple. However, it's at least much better, and perhaps the additional complexity of semantic clauses is simply the price we have to pay when we want compositional semantic values that take into consideration the radically substructural nature of the material inferential relations constitutive of meaning in natural language.

Though these semantic clauses are different than more familiar ones, they do enable us to compositionally compute the semantic values of complex sentences in just the way we should want from a formal semantic framework. Consider, for instance, how this framework can be used to semantically define the notions of contradictions and tautologies, the former being a sentence that cannot be coherently asserted in any position and the latter being a sentence that cannot be denied in any position. Officially:

Contradiction: A is a contradiction just in case $[\![+\langle A\rangle]\!] = \mathbb{P}$

Tautology: A is a tautology just in case $[\![-\langle A\rangle]\!] = \mathbb{P}$

Given the semantic clauses provided, we can show, semantically, that $p \wedge \neg p$ is a contradiction, since every position is incompatible with its assertion:

$$\begin{aligned}
[\![p]\!] &= \langle \{+p\}^{\perp}, \{-p\}^{\perp} \rangle && \mathrm{S}_{\mathcal{A}} \\
[\![\neg p]\!] &= \langle \{-p\}^{\perp}, \{+p\}^{\perp} \rangle && \mathrm{S}_{\neg} \\
[\![+p \wedge \neg p]\!] &= \{\{+p\}^{\perp\perp} \Cup \{-p\}^{\perp\perp}\}^{\perp} && \mathrm{S}_{\wedge} \\
&= \{\{+p\} \cup \{-p\}\}^{\perp} && \text{Fact 2} \\
&= \{+p, -p\}^{\perp} && \text{Set Theory} \\
&= \{\Gamma \mid \Gamma \cup \{+p, -p\} \in \mathbb{I}\} && \text{Definition 4} \\
&= \mathbb{P} && \text{Constraint 2}
\end{aligned}$$

Dually, we can show that $p \vee \neg p$ is a tautology, since for every atomic sentence p, every position containing both $-p$ and $+p$ is in $\mathbb{I}$.

Or consider how we can show that De Morgan's law holds in this framework. We compute the semantic value of $\neg(p \vee q)$ as follows:

$$\begin{aligned}
[\![p]\!] &= \langle \{+p\}^{\perp}, \{-p\}^{\perp} \rangle \\
[\![q]\!] &= \langle \{+q\}^{\perp}, \{-q\}^{\perp} \rangle \\
[\![p \vee q]\!] &= \langle \{+p\}^{\perp} \cap \{+q\}^{\perp}, \{\{-p\}^{\perp\perp} \Cup \{-q\}^{\perp\perp}\}^{\perp} \rangle \\
[\![\neg(p \vee q)]\!] &= \langle \{\{-p\}^{\perp\perp} \Cup \{-q\}^{\perp\perp}\}^{\perp}, \{+p\}^{\perp} \cap \{+q\}^{\perp} \rangle \\
&= \langle \{\{-p\} \cup \{-q\}\}^{\perp}, \{+p\}^{\perp} \cap \{+q\}^{\perp} \rangle \\
&= \langle \{-p, -q\}^{\perp}, \{+p\}^{\perp} \cap \{+q\}^{\perp} \rangle
\end{aligned}$$

We compute the semantic value of $\neg p \wedge \neg q$ as follows, observing that we obtain the same result:

$$
\begin{aligned}
[\![p]\!] &= \langle \{+p\}^{\perp}, \{-p\}^{\perp} \rangle \\
[\![q]\!] &= \langle \{+q\}^{\perp}, \{-q\}^{\perp} \rangle \\
[\![\neg p]\!] &= \langle \{-p\}^{\perp}, \{+p\}^{\perp} \rangle \\
[\![\neg q]\!] &= \langle \{-q\}^{\perp}, \{+q\}^{\perp} \rangle \\
[\![\neg p \wedge \neg q]\!] &= \langle \{\{-p\}^{\perp\perp} \Cup \{-q\}^{\perp\perp}\}^{\perp}, \{+p\}^{\perp} \cap \{+q\}^{\perp} \rangle \\
&= \langle \{\{-p\} \cup \{-q\}\}^{\perp}, \{+p\}^{\perp} \cap \{+q\}^{\perp} \rangle \\
&= \langle \{-p, -q\}^{\perp}, \{+p\}^{\perp} \cap \{+q\}^{\perp} \rangle
\end{aligned}
$$

In this way, we can show, semantically, that $\neg(p \vee q)$ and $\neg p \wedge \neg q$ have the same significance in the sense that they are incoherent to assert in same positions and incoherent to deny in the same positions. In particular, as the final clause states, the set of positions relative to which asserting $\neg p \wedge \neg q$ (or $\neg(p \vee q)$) is incoherent is the same as the set relative to which denying p along with denying q is incoherent, and the set of positions relative to which denying $\neg p \wedge \neg q$ (or $\neg(p \vee q)$) is incoherent is the set relative to which asserting p is incoherent and also asserting q is incoherent.

Finally, let us return to the bilateral reading of *proof-theoretic* consequence, which says that $X \vdash Y$ just in case asserting everything in X and denying everything in Y is incoherent. To consider just the simple case where X and Y are singletons, $A \vdash B$ says that asserting A and denying B is incoherent. We can now define a corresponding notion of *semantic consequence*. As with before, we must define $A \vDash B$ in terms of the *semantic values* of A and B—in terms of the incoherence profiles of A and B, as articulated by our semantic theory. Once again, applying $^{\perp}$ to these semantic values, $[\![A]\!]^{\perp}$ is the set of things that incompatibility entail A. If we consider $[\![+A]\!]^{\perp} \Cup [\![-B]\!]^{\perp}$, then, we have the set of positions consisting in every combination of a position that incompatibility entails $+A$ along with a position that incompatibility entails $-B$. Clearly, insofar as asserting A along with denying B is incoherent, every such position must be incoherent as well. Moreover, every position that incompatibility entails every such position must be incoherent. Kaplan's definition of entailment is that A

entails B just in case this is so. That is, $A \vDash B$ just in case $\{[\![+A]\!]^{\perp} \Cup [\![-B]\!]^{\perp}\}^{\perp\perp} \subseteq \mathbb{I}$. Generalizing:

> **Semantic Entailment:** $A_1, A_2 \ldots A_n \vDash B_1, B_2 \ldots B_n$ just in case $([\![+A_1]\!]^{\perp} \Cup +[\![A_2]\!]^{\perp} \Cup \ldots + [\![A_n]\!]^{\perp} \Cup [\![-B_1]\!]^{\perp} \Cup [\![-B_2]\!]^{\perp} \Cup \ldots [\![-B_n]\!]^{\perp})^{\perp\perp} \subseteq \mathbb{I}$

Kaplan's main result is that given a *base set of incoherent positions* consisting in the assertions and denials of atomics in the semantics, corresponding to a *base consequence relation* between atomics in the proof theory, this semantics is sound and complete with respect to the extension of that base consequence relation yielded by Ketonen's classical sequent calculus: given a base B, $X \vdash_B Y$ just in case $X \vDash_B Y$.[15]

7 Conclusion

I have laid out what I take to be the most philosophically well-motivated interpretation of the semantic framework that Brandom describes as "the current state of the art in inferentialist semantics." Though I have not gone through all the technical details, I hope I have done enough conceptual work that one can now read through those details in Kaplan's dissertation, interpreting them in the way I have laid out. Whether or not this semantic framework is indeed the "grail" that inferentialists have forever been searching for remains to be seen. Still, I hope I've done enough here to show that further developments of it are worth pursuing. The most obvious avenue for further development is to extend implication space semantics, presented here for *sentential* vocabulary, to *subsentential* vocabulary. In forthcoming work, Hlobil [9] has begun to develop it in this direction. Let me conclude by considering one other possible avenue for further development.

I have articulated the "implication space semantics" put forward by Kaplan as a bilateral successor to Brandom's incompatibility semantics which, unlike Brandom's version, is capable of accommodating defeasible material incompatibilities. This is in line with Hlobil and Brandom's [10] official articulation

[15] See Kaplan [13, 234-246] for the details.

of the pragmatic significance of a multiple conclusion "implication," following Restall [20], as telling us that it's incoherent to assert all of the premises and deny all of the conclusions. However, the thought that implication in general really can be reduced to incoherence in this way is quite an un-Brandomian thesis. Even at the time of proposing incompatibility semantics, Brandom explicitly acknowledged that since "incompatibility relations are only *one* dimension of inferential articulation, this semantic representation of conceptual content will necessarily be only partial,"[1, 123 n5]. Though, of course, Brandom did not at the time have *bilateral* incompatibility relations between assertions and denials in view, it's hard to see why this claim wouldn't just as well be applicable here. In *Reason for Logic, Logic for Reasons*, Brandom's official proposal is to treat incoherences of assertions and denials as basic and define a relation of committive consequence on that basis. Concretely, the idea is that if, given one's assertions and denials, it is *incoherent* to assert/deny a sentence A, then one is implicitly *committed* to denying/asserting A. There is substantial reason to think, however, that committive consequence cannot be reduced to incoherence in this way.[16] Consider, for instance, that asserting "*a* is red" along with asserting "*a* is a square" commits one to denying "*a* is blue," and, likewise, asserting "*a* is blue" along with asserting "*a* is a square" commits one to denying "*a* is red." However, it does not seem right to say that asserting "*a* is red" along with asserting "*a* is blue" commits one to denying "*a* is a square." If, however, we want to understand such relations of bilateral committive consequence in terms of bilateral incoherence in the way Brandom suggests, there is no way to draw this distinction. Given the fact that asserting "*a* is red," asserting "*a* is blue," and asserting "*a* is a square" is incoherent, we would, by Brandom's principle, be able to indiscriminately derive all three of these consequence relations.

A natural (and, indeed, Brandomian) response to this issue is to treat both incoherence and committive consequence as equally primitive. In a number of papers [23], [27], [26], [25], I have developed bilateral sequent calculi, of both the single conclusion and multiple conclusion variety, that do just that. These

[16] I give just one example here, which I give in [24] in arguing against the idea that incompatibility between states of affairs can be understood in terms of the incoherence of sets of states of affairs. For different considerations against Brandom's official proposal, see my [27] and Incurvati [11].

calculi have the feature that the *solely left-sided fragment*, translated along the lines mentioned above, correspond to the multiple conclusion sequent calculus, interpreted in Restall-style fashion. However, they also have a bilateral consequence relation *across the turnstile*. If the coordination principles corresponding to Brandom's proposal stated above are not imposed, these two aspects of the sequent calculi (respectively codifying incoherence and committive consequence) are treated as equally basic and irreducible to each other. If, then, a true bilateral implication space semantics could be articulated for such systems (with the incompatibility interpretation laid out above provided a semantics for their solely left-sided fragment), this could provide a single semantics in which both key dimensions of inferential articulation—implication and incompatibility—are truly unified. Doing this, however, presents either a conceptual or a technical challenge. Technically, it is very straightforward to extend this approach to multiple conclusion bilateral sequent calculi.[17] However, one then faces the conceptual challenge of making good sense of multiple conclusion "implications" as *implications* in the proper sense of the term.[18] On the other hand, it is conceptually very straightforward to make sense of single conclusion bilateral sequents as expressing implications, in the proper sense of the term: for instance, as expressing relations of committive or permissive consequence. However, one then faces the technical challenge of providing an incompatibility semantics of the sort presented here for *single conclusion* systems. I'm not sure which route is more promising. In any case, it seems to me that we inferentialists still have a lot of work to do; we have not yet arrived at "the one far-off divine event towards which the whole creation moves." Perhaps, however, we're inching ever closer.

[17]One simply distinguishes not just assertions from denials but also premisory assertions/denials from conclusory assertions/denials, expanding the semantic values from pairs to quadruples in the obvious way.

[18]If one wants to accommodate examples of the sort mentioned above, the doubly-bilateral interpretation of multiple conclusion bilateral sequents I've suggested in some of my work [**?**], where $\Gamma \vdash \Delta$ is read as saying that *making* all the moves in Γ and *challenging* all the moves in Δ, is not available.

References

[1] Robert Brandom. *Between Saying and Doing*. Oxford University Press, Oxford, 2008.

[2] Robert Brandom. Lecture 9. metavocabularies of reason. semantics ii: Implication-space semantics and conceptual roles. Course lecture recording, 2023. `https://www.youtube.com/watch?v=5nDnZjm9o0E`.

[3] Kris Brown. Virtual double categories: A formal model of open reason relations for neurosymbolic ai. Talk presented at the RRR Conference (Reasons, Rationality, and Rationalism), September 2025.

[4] Camillo Fiore. Reading conclusions conjunctively. *Journal of Philosophical Logic*, 53(6):1641–1672, 2024.

[5] Jean-Yves Girard. Linear logic. *Theoretical Computer Science*, 50:1–101, 1987.

[6] Ulf Hlobil. A nonmonotonic sequent calculus for inferentialist expressivists. In Pavel Arazim and Michal Danč'ak, editors, *The Logica Yearbook 2015*, pages 87–105. College Publications, 2016.

[7] Ulf Hlobil. When structural principles hold merely locally. In Pavel Arazim and Tomávs Lávička, editors, *The Logica Yearbook 2016*, pages 53–67. College Publications, 2017.

[8] Ulf Hlobil. Choosing your nonmonotonic logic: A shopper's guide. In Pavel Arazim and Tomáš Lávička, editors, *The Logica Yearbook 2017*, pages 109–123. College Publications, 2018.

[9] Ulf Hlobil. First-order implication-space semantics. *Journal of Philosophical Logic*, Forthcoming.

[10] Ulf Hlobil and Robert Brandom. *Reasons for Logic, Logic for Reasons: Pragmatics, Semantics, and Conceptual Roles*. Routledge, New York, 2024.

[11] Luca Incurvati. A plea for multilateralism. *Philosophical Studies*, Forthcoming.

[12] Daniel Kaplan. A multi-succident sequent calculus for logical expressivists. In Pavel Arazim and Tomáš Lákiva, editors, *The Logica Yearbook 2017*, pages 139–154. College Publications, 2018.

[13] Daniel Kaplan. *Substructural Content*. PhD thesis, University of Pittsburgh, 2022.

[14] Oiva Ketonen. Untersuchungen zum prädikatenkalkül. *Annales Academiae Scientiarum Fennicae Series A, I. Mathematica-physica*, 1944.

[15] Mark Lance. The logical structure of linguistic commitment iii brandomian scorekeeping and incompatibility. *Journal of Philosophical Logic*, 30(5):439–464, 2001.

[16] Mark Lance and Philip Kremer. The logical structure of linguistic commitment i: Four systems of non-relevant commitment entailment. *Journal of Philosophical Logic*, 23(4):369–400, 1994.

[17] Mark Lance and Philip Kremer. The logical structure of linguistic commitment ii: Systems of relevant commitment entailment. *Journal of Philosophical Logic*, 25(4):425–449, 1996.

[18] Bernhard Nickel. Dynamics, brandom-style. *Philosophical Studies*, 162(2):333–354, 2013.

[19] Daniele Porello. Incompatibility semantics from agreement. *Philosophia*, 40(1):99–119, 2012.

[20] Greg Restall. Multiple conclusions. In Petr Hájek, Luis Valdés-Villanueva, and Dag Westerståhl, editors, *Logic, Methodology and Philosophy of Science*. College Publications, 2005.

[21] Ellie Ripley. Paradoxes and failures of cut. *Australasian Journal of Philosophy*, 91(1):139–164, 2013.

[22] Ryan Simonelli. Considering the exceptions: On the failure of cumulative transitivity for indicative conditionals. *Synthese*, 200(5):1–21, 2022.

[23] Ryan Simonelli. A general schema for bilateral proof rules. *Journal of Philosophical Logic*, (3):1–34, 2024.

[24] Ryan Simonelli. Why must incompatibility be symmetric? *Philosophical Quarterly*, 74(2):658–682, 2024.

[25] Ryan Simonelli. There is a logical negation: ″yes,″ ″no,″ both, neither. *Australasian Journal of Logic*, 23(1):55–78, 2026.

[26] Ryan Simonelli. “yes, “no,” neither, and both. *Synthese*, 207(1):34, 2026.

[27] Ryan Simonelli. Brining bilateralisms together: A unified framework for inferentialists. *Australasian Journal of Philosophy*, Forthcoming.

[28] Florian Steinberger. Why conclusions should remain single. *Journal of Philosophical Logic*, 40(3):333–355, 2011.

Appendix: Semantic Values, Hlobil-Style

The semantic values I have presented in this paper are those presented in Kaplan’s dissertation. As I mentioned in the introduction, however, these semantic values are not the ones put forward in Hlobil and Brandom’s *Reasons* *for Logic, Logic for Reasons*; Hlobil's approach there is to lift semantic values from sets of positions to sets of sets of positions. I mentioned that comprehending what the semantic value of a simple sentence such as "*a* is red and round" actually is, on this approach, can feel like an impossible task. As a service to readers of Hlobil and Brandom's book, this appendix undertakes this task with the aid of the conceptual resources developed here.

Rather than taking the semantic value of a position to be its incoherence profile, Hlobil takes it to be the equivalence class of sets of positions that have that profile. Thus, in this formulation, rather than taking the semantic value of a position Γ to be $\Gamma^{\perp}$, we take it to be $\{\mathbf{X} \mid \mathbf{X}^{\perp} = \Gamma^{\perp}\}$. Thus, for instance, the semantic value of $\{+\mathbf{red}(a)\}$ is not the incoherence profile of $\{+\mathbf{red}(a)\}$ (the set of all positions incompatible with $\{+\mathbf{red}(a)\}$) but, rather, the set of all sets of positions that have this incoherence profile. What are these sets of positions? Of course, one such set is $\{\{+\mathbf{red}(a)\}\}$, as well as any set containing only positions that have the same incoherence profile as $\{+\mathbf{red}(a)\}$, for instance, $\{\{+\mathbf{red}\}, \{+\mathbf{primary\ color}, -\mathbf{yellow}, -\mathbf{blue}\}\}$. But it will also contain any set that adds to such a set any position whose incoherence profile is strictly stronger than $\{+\mathbf{red}(a)\}$. For instance, it will include $\{\{+\mathbf{red}\}, \{+\mathbf{scarlet}\}\}$. This is because the incoherence profile of a set of positions is the intersection of the incoherence profiles of its members, and so the incoherence profile of such a set will still be the same as $\{+\mathbf{red}(a)\}$. So, on Hlobil's approach, the semantic value of a position Γ is the set of sets of positions such that each element in each of those sets incompatibility entails Γ and also includes at least one member incompatibility equivalent to Γ. That is, $[\![\Gamma]\!]$ is the subset of $\mathcal{P}(\Gamma^{\perp\perp})$ whose members each contain at least one element Δ such that $\Delta^{\perp} = \Gamma^{\perp}$.

Hlobil speaks of such equivalence classes as the "role" of a set of positions. Officially, the role of a set of positions $\mathbf{X}$, $\mathcal{R}(\mathbf{X})$, is the set containing the sets of positions that each have the same incoherence profile as $\mathbf{X}$. They're possessed by individual positions in the derivative sense that the singleton containing that position has a role. Officially:

Definition 6.1 (Roles):

6.1a: For any $\mathbf{X} \subseteq \mathbb{P}$, $\mathcal{R}(\mathbf{X}) = \{\mathbf{Y} \mid \mathbf{Y}^{\perp} = \mathbf{X}^{\perp}\}$

6.1b: For any $\Gamma \in \mathbb{P}$, $\mathcal{R}(\Gamma) = \mathcal{R}(\{\Gamma\}) = \{\mathbf{X} \mid \mathbf{X}^{\perp} = \{\Gamma\}^{\perp}\}$

6.1c: For any $A \in \mathcal{L}$, $\mathcal{R}(A) = \langle \mathcal{R}(\{+A\}), \mathcal{R}(\{-A\}) \rangle$

Generalizing the above points just informally indicated about the "role" of an individual position, the following facts hold about the "roles" of sets of positions:

Fact 3: $\mathbf{X} \in \mathcal{R}(\mathbf{X})$

Fact 4: $\mathbf{X}^{\perp\perp} \in \mathcal{R}(\mathbf{X})$

Fact 5: $\mathcal{R}(\mathbf{X}) \subseteq \mathcal{P}(\mathbf{X}^{\perp\perp})$

Hlobil defines two operations that combine roles in different ways:

Adjunction: $\mathcal{R}(\mathbf{X}) \sqcup \mathcal{R}(\mathbf{Y}) = \mathcal{R}(\mathbf{X} \Cup \mathbf{Y})$

Symjunction: $\mathcal{R}(\mathbf{X}) \sqcap \mathcal{R}(\mathbf{Y}) = \mathcal{R}(\mathbf{X} \cup \mathbf{Y})$

To adjoin $\mathcal{R}(\mathbf{X})$ with $\mathcal{R}(\mathbf{Y})$, one takes the set consisting in all of the unions of an element of **X** along with an element of **Y**, and then takes the role of that set. To symjoin $\mathcal{R}(\mathbf{X})$ with $\mathcal{R}(\mathbf{Y})$, one unions **X** with **Y**, and then takes the role of that set. Considering just these operations on the roles of a pair of positions, adjoining the roles of two positions is taking the set of positions that have the incoherence profile of the position consisting in the union of those two positions, whereas symjoining the roles of two positions is taking the set of positions that have the incoherence profile that is the intersection of the incoherence profiles of those two positions.

With these two operations defined in this way, Hlobil defines the semantic clauses for logically complex sentences as follows:

Semantic Clauses: $[\![A]\!] = \langle [\![+A]\!], [\![-A]\!] \rangle$ where:

$\mathbf{S}_{\mathcal{A}}$: $[\![p]\!] = \langle \mathcal{R}(\{\{+p\}\}), \mathcal{R}(\{\{-p\}\}) \rangle$

$\mathbf{S}_{\neg}$: $[\![\neg A]\!] = \langle [\![-A]\!], [\![+A]\!] \rangle$

$\mathbf{S}_{\wedge}$: $[\![A \wedge B]\!] = \langle [\![+A]\!] \sqcup [\![+B]\!], [\![-A]\!] \sqcap [\![-B]\!] \rangle$

$\mathbf{S}_{\vee}$: $[\![A \vee B]\!] = \langle [\![+A]\!] \sqcap [\![+B]\!], [\![-A]\!] \sqcup [\![-B]\!] \rangle$

$\mathbf{S}_{\rightarrow}$: $[\![A \rightarrow B]\!] = \langle [\![-A]\!] \sqcap [\![+B]\!], [\![+A]\!] \sqcup [\![-B]\!] \rangle$

To see how these clauses are meant to work, consider the derivation of the equivalence of $\neg(p \vee q)$ and $\neg p \wedge \neg q$ in this framework:

$$
\begin{aligned}
[\![p]\!] &= \langle \mathcal{R}(\{\{+p\}\}), \mathcal{R}(\{\{-p\}\})\rangle \\
[\![q]\!] &= \langle \mathcal{R}(\{\{+q\}\}), \mathcal{R}(\{\{-q\}\})\rangle \\
[\![p \vee q]\!] &= \langle \mathcal{R}(\{\{+p\}\}) \sqcap \mathcal{R}(\{\{+q\}\}), \mathcal{R}(\{\{-p\}\}) \sqcup \mathcal{R}(\{\{-q\}\})\rangle \\
[\![\neg(p \vee q)]\!] &= \langle \mathcal{R}(\{\{-p\}\}) \sqcup \mathcal{R}(\{\{-q\}\}), \mathcal{R}(\{\{+p\}\}) \sqcap \mathcal{R}(\{\{+q\}\})\rangle \\
&= \mathcal{R}(\{\{-p, -q\}\}), \mathcal{R}(\{\{+p\}, \{+q\}\})
\end{aligned}
$$

$$
\begin{aligned}
[\![p]\!] &= \langle \mathcal{R}(\{\{+p\}\}), \mathcal{R}(\{\{-p\}\})\rangle \\
[\![q]\!] &= \langle \mathcal{R}(\{\{+q\}\}), \mathcal{R}(\{\{-q\}\})\rangle \\
[\![\neg p]\!] &= \langle \mathcal{R}(\{\{-p\}\}), \mathcal{R}(\{\{+p\}\})\rangle \\
[\![\neg q]\!] &= \langle \mathcal{R}(\{\{-q\}\}), \mathcal{R}(\{\{+q\}\})\rangle \\
[\![\neg p \wedge \neg q]\!] &= \langle \mathcal{R}(\{\{-p\}\}) \sqcup \mathcal{R}(\{\{-q\}\}), \mathcal{R}(\{\{+p\}\}) \sqcap \mathcal{R}(\{\{+q\}\})\rangle \\
&= \langle \mathcal{R}(\{\{-p, -q\}\}), \mathcal{R}(\{\{+p\}, \{+q\}\})\rangle
\end{aligned}
$$

There is, however, a technical issue: adjunction and symjunction, as defined above, are defined *syntactically*; they are not really defined as operations on the set-theoretic objects that are the roles themselves. That is, once we actually *take* the roles $\mathbf{X}$ and $\mathbf{Y}$, we are left with two set-theoretic objects (sets of sets of positions), and the official definition Hlobil provides does not tell us how to operate on these set-theoretic objects that are the roles themselves in order to arrive at the symjunctions and adjunctions of roles that are the candidate semantic values for logically complex sentences. Accordingly, the semantic clauses stated above do not technically meet the compositionality requirement. However, it is easy to modify the definitions of adjunction and symjunction so that they do.

An important feature of $\mathcal{R}(\mathbf{X}) \sqcup \mathcal{R}(\mathbf{Y})$ (and $\mathcal{R}(\mathbf{X}) \sqcap \mathcal{R}(\mathbf{Y})$) is that it actually doesn't matter what elements of $\mathcal{R}(\mathbf{X})$ and $\mathcal{R}(\mathbf{Y})$ you select, union (or pairwise union), and then take the role of.[19] You could take **X** and **Y** (which, recall are elements of $\mathcal{R}(\mathbf{X})$ and $\mathcal{R}(\mathbf{Y})$) as the official definition of the adjunction of roles suggests, or you could take any random elements of **X** and **Y** you want. Thus, where $\mathsf{pick}(x)$ denotes the result of picking an arbitrary element of x, Adjunction and Symjunction can just as well be defined as follows:

Adjunction: $\mathcal{R}(\mathbf{X}) \sqcup \mathcal{R}(\mathbf{Y}) = \mathcal{R}(\mathsf{pick}(\mathcal{R}(\mathbf{X})) \Cup \mathsf{pick}(\mathcal{R}(\mathbf{Y})))$

Symjunction: $\mathcal{R}(\mathbf{X}) \sqcap \mathcal{R}(\mathbf{Y}) = \mathcal{R}(\mathsf{pick}(\mathcal{R}(\mathbf{X})) \cup \mathsf{pick}(\mathcal{R}(\mathbf{Y})))$

Defining adjunction and symjunction in this way, they really are defined as an operation on the roles. This is important for the clauses stated by Hlobil actually meeting the compositionality requirement. Of course, given that we can pick *any* member of these roles, we can also define these operations in such a way that we pick *specific* member(s) of these roles. One special member of $\mathcal{R}(\mathbf{X})$, for instance, is $\mathbf{X}^{\perp\perp}$. It follows directly from facts 4 and 5 stated above that this special member of $\mathcal{R}(\mathbf{X})$ is identical to the *union* of all of the members of $\mathcal{R}(\mathbf{X})$. Thus, we could just as well define adjunction and symjunction as follows:

Adjunction: $\mathcal{R}(\mathbf{X}) \sqcup \mathcal{R}(\mathbf{Y}) = \mathcal{R}(\bigcup(\mathcal{R}(\mathbf{X})) \Cup \bigcup(\mathcal{R}(\mathbf{Y})))$

Symjunction: $\mathcal{R}(\mathbf{X}) \sqcap \mathcal{R}(\mathbf{Y}) = \mathcal{R}(\bigcup(\mathcal{R}(\mathbf{X})) \cup \bigcup(\mathcal{R}(\mathbf{Y})))$

$\mathbf{X}^{\perp\perp}$ is the *biggest* element of $\mathcal{R}(\mathbf{X})$ in that every other element of $\mathcal{R}(\mathbf{X})$ is a subset of $\mathbf{X}^{\perp\perp}$. On the flip side, there is generally no unique *smallest* element of $\mathcal{R}(\mathbf{X})$ of which every other element of $\mathcal{R}(\mathbf{X})$ is a superset. However, where **X** is itself a singleton, as it will be when we take the semantic values of sentences, it is nevertheless fruitful to consider the set of singletons in $\mathcal{R}(\mathbf{X})$. Where we have $\mathcal{R}(\{\Gamma\})$ for some position Γ the set of singletons will, of course, include $\{\Gamma\}$, but it will also include any position that is incompatibility-equivalent to Γ. For instance, the set of singletons in $\mathcal{R}(\{\{+\mathbf{red}(a)\}\})$ includes not only $\{\{+\mathbf{red}(a)\}$ but also $\{\{+\mathbf{primary\ color}(a), -\mathbf{blue}(a), -\mathbf{yellow}(a)\}\})$. Thus, where $\mathsf{min}(x)$ denotes the subset of x whose elements have the least members, we may also define adjunction and symjunction as follows:

[19] See Hlobil and Brandom [10, 256-257] for proofs.

Adjunction: $\mathcal{R}(\mathbf{X}) \sqcup \mathcal{R}(\mathbf{Y}) = \mathcal{R}(\mathsf{pick}(\mathsf{min}(\mathcal{R}(\mathbf{X}))) \Cup \mathsf{pick}(\mathsf{min}(\mathcal{R}(\mathbf{Y}))))$

Symjunction: $\mathcal{R}(\mathbf{X}) \sqcap \mathcal{R}(\mathbf{Y}) = \mathcal{R}(\mathsf{pick}(\mathsf{min}(\mathcal{R}(\mathbf{X}))) \cup \mathsf{pick}(\mathsf{min}(\mathcal{R}(\mathbf{Y}))))$

Insofar as all of the entities that are symjoined or adjoined are roles, all of these definitions define precisely the same operation. However, when it comes to their application in defining semantic values, this is perhaps the closest definition to Hlobil's official definition of these operations. Unlike Hlobil's official definition, however, this actually defines adjunction and symjunction as binary operations on the set-theoretic entities that are the roles themselves. Instead of taking the role of the (pairwise or regular) union of two positions Γ and Δ, we take the role of two positions that are incompatibility equivalent with Γ and Δ.[20]

To take a concrete example, consider the semantic value of "*a* is red and *a* is round." The semantic value of "*a* is red" is the pair consisting in the role of its assertion along with the role of its denial. Let us consider each of the elements of this pair in turn. The role of asserting "*a* is red," $\mathcal{R}(\{\{+\mathbf{red}(a)\}\})$, is the set of sets of positions such that each element in each such set incompatibility entails asserting "*a* is red" and each such set includes at least one position that is incompatibility-equivalent to asserting "*a* is red." Similarly, $\mathcal{R}(\{\{+\mathbf{round}(a)\}\})$ is such a set. Now, the role of asserting "*a* is red and round," Hlobil's semantics tells us, is the adjunction of these two sets. Using the third definition of adjunction corresponds most closely to Hlobil's definition that actually lets us compute this role, let us take a singleton $\{\Gamma\}$ in $\mathcal{R}(\{\{+\mathbf{red}(a)\}\})$ and a singleton $\{\Delta\}$ in $\mathcal{R}(\{\{+\mathbf{round}(a)\}\})$, pairwise union these sets to obtain $\{\Gamma \cup \Delta\}$, and then take the role of this set.[21] Given that Γ is incompatibility-equivalent with $\{+\mathbf{red}(a)\}$ and Δ is incompatibility equivalent with $\{+\mathbf{round}(a)\}$, this role is the same as that of $\{\{+\mathbf{red}(a), +\mathbf{round}(a)\}\}$. Thus, the role of asserting "*a* is red and *a* is round" is identical to the role of asserting "*a* is red" along with asserting "*a* is round."

Consider now the role of denying "*a* is red and *a* is round." This is the

[20] As far as I can tell, there's no pure set-theoretic operation that enables us to recover Γ and Δ themselves, once we take their roles.

[21] I leave it as an exercise for the reader to think through how the computation works with the first and the second definition.

symjunction of the role of denying "*a* is red" and the role of denying "*a* is round." Once again, the role of denying "*a* is red" is the set of sets of positions such that each position in each such set incompatibility entails denying "*a* is red" and each such set includes at least one position that is incompatibility equivalent to denying "*a* is red." Similarly for the role of denying "*a* is round." In keeping with the previous approach, let us take the symjunction of these two sets by taking a singleton $\{\Gamma\}$ belonging to the first set, which will have the same incoherence profile as denying "*a* is red," and a singleton $\{\Delta\}$ belonging to the second set, which will have same the incoherence profile as denying "*a* is round." The symjunction of $\mathcal{R}(\{\{-\mathbf{red}(a)\}\})$ and $\mathcal{R}(\{\{-\mathbf{round}(a)\}\})$ will then be the role of the union $\{\Gamma\}$ and $\{\Delta\}$. That is, it's the role of $\{\Gamma, \Delta\}$. This is the set of sets of positions whose incoherence profile is the same as $\{\Gamma, \Delta\}$. Now, once again, the incoherence profile of a set of positions is the intersection of the incoherence profiles of its members. Given the equivalence of Γ with denying "*a* is red" and Δ with denying "*a* is round," the incoherence profile of $\{\Gamma, \Delta\}$ will be the set of positions such that, occupying them, denying "*a* is red" is incoherent and also denying "*a* is round" is incoherent. The *role* of this set, then, will be the set of sets of positions that have the same incoherence profile as this set. The inferentially *weakest* member of each such set will be some position that is *minimally incompatible* both with denying "*a* is red" and with denying "*a* is round." Presumably, this will be the position consisting of asserting "*a* is red" along with asserting "*a* is round" or any position that is incompatibility-equivalent to this position. So, the role of $\{\{-\mathbf{red}(a)\}, \{-\mathbf{round}(a)\}\}$ will be the set of sets of positions such that each position in each such set incompatibility entails $\{+\mathbf{red}(a), +\mathbf{round}(a)\}$ (for instance, $\{+\mathbf{scarlet}(a), +\mathbf{circle}(a)\}$) and each set includes at least one position that is incompatibility equivalent to $\{+\mathbf{red}(a), +\mathbf{round}(a)\}$.

I have gone through the pains of actually explaining Hlobil's version of implication space semantics primarily in service of (what I expect to be many) frustrated readers of *Reasons for Logic, Logic for Reasons* reaching the final substantive chapter only to find its content impenetrable. In doing this, I take myself to have shown at least two things. First, the semantic values defined in this version of the framework *are* in fact intelligible despite the strain on cognitive resources that comprehending them actually takes. Second, despite the semantic clauses put forward by Kaplan looking more complex than those put forward by Hlobil, the semantic values defined in the version of the framework developed by Kaplan are in fact much easier to comprehend than those defined by Hlobil. Moreover, once we actually formulate clauses in Hlobil's variant that meet the compositionality requirement, they do not look much simpler than Kaplan's at all. To what extent it makes sense to ascend to the level of abstraction at which Hlobil's version of the semantic framework operates will presumably depend on one's purposes. For certain technical purposes, ascending to the level of Hlobil's "roles" may be convenient. However, in terms of comprehending the core philosophical idea of the framework (and I suspect for most technical applications as well), it seems to me that Kaplan's semantic values are preferable.