2025 Samuel Pufendorf Lectures
Lund University
May 20, 21, and 22, 2025

Vocabularies of Reason

Bob Brandom

I Reasoning and Representing 8

II Logic and the Structure of Reasons 41

III Roles and Reasons 75

Abstract of the 2025 Samuel Pufendorf Lectures:

Vocabularies of Reason

What is it to live, as we do, in the space of reasons? Reporting work from a recent book, these lectures explore how reasons show up from the perspectives afforded by different ways of talking about them. The emphasis throughout is on the *relational* character of reasons: in particular, the relations of implication and incompatibility. These show up as articulating essential norms governing the *use* of declarative sentences to make claims, underwriting practical assessments of rational defenses of claims, by giving reasons for them, and rational challenges to claims, by giving reasons against them. The same reason relations show up in a different guise in truth-maker theories of the *meaning* of declarative sentences. In addition to these pragmatic and semantic vocabularies for talking about reason relations, I consider logical vocabularies for making reason relations explicit (recommending one as the best at that), and introduce a new formal language for talking about and manipulating the conceptual roles sentences come to play by standing in reason relations. Considering the relations among these perspectives on reasons yields a sketch of a higher form of rational self-consciousness.

Preamble:

In these three lectures I present a narrative path through some of the notable features of the recently published book *Reasons for Logic, Logic for Reasons: Pragmatics, Semantics, and Conceptual Roles*.[1] That volume is the result of intensive work over the past decade by the logic group **Research on Logical Expressivism** (ROLE). The principal author of the book is Ulf Hlobil; I am his co-author. The origins of some of our crucial technical results lie in the 2021 University of Pittsburgh Ph.D. dissertation of ROLE-contributor Daniel Scott Kaplan. Other important arguments and considerations are due to the work of long-time ROLE-contributors Ryan Simonelli, Rea Golan, and Shuhei Shimamura. The amount and complexity of the material I am offering an overview of in these lectures means that my summaries must sometimes omit details and qualifications, as well as demonstrations, that can be found in the original.

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

Abstract of Lecture 1:

Reasoning and Representing

This lecture introduces and develops the idea that the crucial representational relation between discursive practices and the world they make it possible for their participants to make claims *about* can be understood to begin with at the level of relations of consequence and incompatibility, rather than at the level of using terms and predicates referring to objects and relations, or even at the level of the facts that the basic uses of declarative sentences purport to state. In a bilateral model of a minimal linguistic practice, relations of implication and incompatibility show up in normative form, as some constellations of commitments to assert and deny sentences preclude entitlement to others. In a truthmaker semantic framework for specifying how the world must be for various sentences to be true or false, relations of consequence and incompatibility show up in the alethic modal metaphysical form of the impossibility of the states resulting from mereologically fusing various sets of truth-making and falsity-making states. These accounts can be arranged so as to be systematically isomorphic—to share a conceptual structure articulated by those reason relations. Such a view has illuminating historical antecedents.

Pufendorf Lectures 2025: **Vocabularies of Reason**

Handout for Lecture 1

Reasoning and Representing

Bob Brandom

Full Disclosure: These lectures are shamelessly promoting some of the ideas developed in more detail in the book: Ulf Hlobil and Robert Brandom, *Reasons for Logic, Logic for Reasons: Pragmatics, Semantics, and Conceptual Roles*, [Routledge, 2024].

Pragmatics must answer to semantics: It is a criterion of adequacy on an account of the *use* of language that it show how such use serves to confer *meaning* on linguistic expressions—establishing the association of some sort of propositional contents (as construed by the semantics), with sentences of the language.
Semantics must answer to pragmatics: It is a criterion of adequacy on an account of the *meanings* of linguistic expressions that it show how meanings codify or determine norms governing the *use* of those expressions—paradigmatically, the use of declarative sentences to make claims or assertions.

Gilbert Harman: "There is no such thing as deductive inference." We must distinguish *relations* of implication from inferential *practices*. (1984) "Logic and reasoning." Synthese, 60(1):107–127.

A minimal model of discursive practice includes:

- Speech acts of *assertion* and *denial* by uttering declarative sentences, which express
- Practical attitudes of *accepting* and *rejecting*, which are
- Doxastic *commitments*, *entitlement* to which can be
- *Challenged* and *defended* by further claims.
- Reason relations of *implication* (consequence) and *incompatibility*, determining which claimables are reasons *for* and reasons *against* other claimables, and so which are suitable as defenses and challenges.

Greg Restall's and David Ripley's bilateral normative pragmatics for the sequent calculus defines reason relations:

- Γ *implies* A iff the position of being committed to *accept* all of Γ and to *reject* A is "out of bounds": a constellation of commitments to which one cannot be entitled.

- Γ is *incompatible* with A iff the position of being committed to *accept* all of Γ and to *accept* A is "out of bounds": a constellation of commitments to which one cannot be entitled.
- If *commitment* to *accept* all of Γ precludes *entitlement* to *reject* A, then it *implicitly* commits one to *accept* A.

Kit Fine's truth-maker semantics and its modal, mereological metaphysics:

- A universe of states,
- Divided into possible and impossible states. (Modal structure)
- States can be fused with others to form new states as wholes, of which they are parts. (Mereological structure)
- A semantic interpretation function assigns declarative sentences to pairs of sets of states, understood as the truth-makers and falsity-makers (verifiers and falsifiers) of those sentences, subject to the condition of
- Exclusivity: every fusion of truth-makers of a sentence with any falsity-maker of that sentence is an impossible state.
- Consequence as Entailment: Γ *entails* A iff every verifier of all of Γ is a verifier of A.
- Consequence as Containment: A *contains* Γ iff every verifier of A includes as a part a verifier of all of Γ and every verifier of all of Γ is a part of a verifier of A.
- There are many more propositions ($=_{df.}$ pairs of sets of states satisfying Exclusivity) than can be expressed by the sentences of any particular language.

Suggestion: Define *implication* in the truth-maker framework by analogy to Exclusivity.
Definition: Γ *implies* A iff every fusion of any truth-maker of all of Γ with any falsity-maker of A is an impossible state.

Hlobil isomorphism of bilateral normative pragmatic definition of reason relations and truth-maker semantic definition:

i) Pragmatic consequence: Γ implies A iff any position that includes accepting all of Γ and rejecting A is normatively incoherent or "out of bounds": one cannot be entitled to such a constellation of commitments.

ii) Semantic consequence: Γ implies A iff any fusion of a state that verifies all the members of Γ with a state that falsifies A is an impossible state.

iii) Pragmatic incompatibility: Γ is incompatible with A ⇔ the position resulting from concomitant commitment to *accept* all of Γ and to *accept* A is normatively *incoherent* ("out of bounds"): a constellation of commitments to which one *cannot* be entitled.

iv) Semantic incompatibility: Γ is incompatible with A ⇔ the state resulting from *fusion* of any *verifiers* of all the members of Γ with any *verifier* of A is an *impossible* state.

The key to *conceptual realism* is understanding conceptual form as role with respect to reason *relations*, rather than reason*ing*: the Harman distinction. For relations of a kind of *in*clusion (consequence, implication) and *ex*clusion (inconsistency, incompatibility) characterize both discursive thought and the world thought about. We understand these as *reason* relations (so as functionally defining specifically *conceptual* form) because of the role they play pragmatically as norms governing what claimables are properly treated as reasons for and against what other claimables. But it turns out that we can then see relations articulating the world that is there to be represented semantically as *isomorphic* to those that articulate discursive practice.

Reason relations are modally robust. The modality is different on the subjective side of appearance than on the objective side of reality. The modality characteristic of consequence and incompatibility on the pragmatic side of reason relations implicit in practices of rationally defending and challenging (giving reasons for and against) claimings is a *deontic normative* matter of what constellations of *commitments* one can be or is precluded from being *entitled* to. The modality characteristic of consequence and incompatibility on the side of the metaphysics of representational semantics is an *alethic modal* matter of what combinations of states are *possible* or *im*possible. This is *bimodal* conceptual realism. In the third lecture, we will consider the abstract *ranges of subjunctive robustness* of reason relations, which manifest themselves concretely in these two modalities.

1. A key element of early modern philosophers' response to the rise of the new science was to move from thinking of appearance in terms of its *resemblance to* reality to thinking of it in terms of its *representation of* reality.

2. Looking at Descartes' algebraic representation of geometrical properties, Spinoza understood the new notion of representation in holistic terms of a global isomorphism: "the order and connection of ideas is the same as the order and connection of things." (*Ethics* II, Prop 7.)

3. Kant took a further step away from the original *perceptual* paradigm of the appearance/reality distinction by focusing on specifically *conceptual* appearances.

"Kant was on the right track when he insisted that just as concepts are essentially (and not accidentally) items which can occur in judgments, so judgments (and, therefore, indirectly concepts) are essentially (and not accidentally) items which can occur in reasonings or arguments." Wilfrid Sellars "Inference and Meaning" [I-4].

May 20, 2025

Vocabularies of Reason
2025 Pufendorf Lectures
Lund University

Lecture I

Reasoning and Representing

Bob Brandom

I. Semantics and Pragmatics

When I say that cats are mammals, we can distinguish between what I have *said*—something sayable, namely that cats are mammals—and my *saying* of it, which is something I have *done*, a speech act I have performed. Of course this 'ing'/'ed' structure of doing and thing done, act and (direct) object acted on, is a common linguistic construction. We distinguish breakings from what is broken, sortings from what is sorted, and drinkings from what is drunk. In the special case of discursive doings, such as sayings and thinkings, the object said or thought is of a special form or kind. What can in the central sense be said or thought is conceptual, more specifically propositional, *contents*.[2] Such contents are the subject of a special intellectual discipline: semantics. The semantic tradition that runs from Frege through the early Wittgenstein, Carnap, Tarski, and Quine, to Kripke, David Lewis, and Kit Fine among many others deploys a constellation of expressively powerful technical vocabularies for specifying, relating, and combining various kinds of meanings or conceptual contents. It addresses, to begin with, the meanings of logical and mathematical locutions, and those of other artificial languages. But the tradition extends from there, both actually and aspirationally, to encompass the contents of a wide range of kinds of linguistic expression found in natural languages. We can reasonably hope to deploy the metaconceptual resources of semantic theory to explain the difference between discursive contents, which can be said, thought, or meant, and things that can merely be broken, sorted, drunk, or the like.

[2] I am using 'discursive' in Kant's sense, as meaning "of or pertaining to concept-use."

What about the practical, 'ing' side of discursive doings? This is the 'force' side of Frege's force/content distinction, whose paradigm is the assertional force distinctive of saying something in the sense of making a statement or claim. Broadening traditional usage a bit, we can call the study of this practical dimension of discursive activity 'pragmatics'. The idea is that pragmatics studies the *use* of linguistic expressions, and semantics studies their *meanings* (to put the distinction in Wittgensteinian terms).

Rough as it is, this characterization already invites us to think about the obligations each of these theoretical enterprises, semantic and pragmatic, owes to each other. In saying things, we are *doing* something sufficiently different from nondiscursive doings such as breaking, dropping, and drinking that what is said admits of a specifically *semantic* interpretation, as having a *meaning* or conceptual, specifically propositional, *content* in a sense in which what is broken, dropped, or drunk does not. A minimal criterion of adequacy on a pragmatic theory is that it account for this difference. What is it about discursive practice that establishes the crucial connection between assertional speech acts and the propositional contents they express? There is nothing except the *use* of declarative sentences to confer propositional content on them. The semantic interpretability of sentences must be intelligible in broadly functional terms of the roles such linguistic expressions play in discursive practice. Along this dimension, pragmatics must answer to semantics. For an adequate pragmatic theory must underwrite an account of how the relation between sentential expressions and the semantic interpretants assigned to them by some semantic theory is fixed. I will call this the "**conferral**" condition: pragmatics owes an account of how *use*, however it construes use, confers *meaning*, as understood by some semantic theory.

Conversely, there is also a sense in which semantics answers to pragmatics. Formal semantics can be thought of as having only the restricted task of showing how to compute the semantic interpretants of complex expressions based on stipulated associations of semantic interpretants with simple expressions. But the *point* of associating sayable meanings or semantic interpretants with expressions at all is ultimately to make sense of what speakers are *doing* in using those expressions. Appeal to the meanings or contents of speech acts must help explain what practitioners are *doing* in saying something by producing those discursive performances. I

will call this the "**codification**" condition: semantics owes an account of how the *meanings* of expressions explain or constrain their *use*.

One model of how this condition might be satisfied understands the features of use to be explained by associating semantic interpretants with expressions as *dispositions of* and *behavioral regularities exhibited by* language users, described in a spare, naturalistic vocabulary. Meanings are then thought of as theoretical entities, postulated to explain observable patterns of discursive behavior. From this point of view, Dummett's insistence that the co-ordinate character of the concepts meaning and understanding entails that semantic properties must be analyzable without remainder into features of use shows up as a kind of semantic instrumentalism. For it in effect rejects postulating theoretical entities in explaining linguistic behavior. (That is the criticism Sellars makes of Rylean behaviorist rejection of the postulation of thoughts and sensations, in the philosophy of mind.) And Wittgenstein's wholesale rejection of the concept of *meaning* in favor of *use* shows up as a consequence of his conviction that "philosophy is not one of the natural sciences," once the postulation of unobservables to explain the antics of observables is identified as the core of natural scientific methodology.

I do not think this model will help us understand what is distinctive about the specifically discursive form of the act/content distinction: what distinguishes sayables and thinkables from breakables, sortables, and drinkables. For that we must think of meanings as codifying *norms* governing the use of expressions that have those meanings. We should understand what the semantic theorist is doing in associating a certain meaning or semantic interpretant with an expression as undertaking a commitment regarding how it would be *correct* or *proper* to use the expression in question. The immediate consequence of specifying the truth-conditions of a declarative sentence, for instance is to set a standard of appraisal or assessment of the correctness or propriety of assertions of that sentence, in a distinctively semantic sense of 'correct' or 'appropriate.' The usefulness of semantic characterizations of discursive performances in explaining dispositions or regularities of behavior of users of that sentence is indirect, mediated by the use-governing norms those semantic attributions codify.

I have gestured in the direction of an important dimension along which *pragmatics answers to semantics*. It is a criterion of adequacy on an account of the *use* of language that it show how such use serves to confer *meaning* on linguistic expressions—establishing the association of some sort of propositional contents (as construed by the semantics), with sentences of the language. I have also gestured in the direction of an important dimension along which *semantics answers to pragmatics*. It is a criterion of adequacy on an account of the *meanings* of linguistic expressions that it show how meanings codify or determine norms governing the *use* of those expressions—paradigmatically, the use of declarative sentences to make claims or assertions. Because we have so much better worked-out formal *semantic* theories of meaning than we do metasemantic, *pragmatic* theories of use, it is worth thinking of the conferral condition and the codification condition as criteria of adequacy of adequacy for pragmatics (in the very broad sense in which I am using the term). Together, they require that an adequate pragmatic theory must characterize the use of language in such a way as to show how, so specified, discursive practices are intelligible as conferring meanings or conceptual contents on linguistic expressions (paradigmatically, propositional content on declarative sentences), such that those contents or meanings are intelligible as setting normative standards for appraisal of the correctness of linguistic moves.

In the abstract, it is difficult to say what a pragmatics that met those conditions might look like. So, I am going to start by sketching a minimal discursive practice, specified in a simple, regimented pragmatic metavocabulary. Then I will introduce Kit Fine's truthmaker semantic theory and show how it can be understood to be so related to the pragmatic theory that the twin requirements are satisfied. What practitioners *do* in talking, according to the pragmatic theory, both suffices to connect the expressions used to Finean semantic interpretants, and explains how those interpretants set normative standards for assessment of the correctness of speech acts that use those expressions. This is the first large move in a narrative arc will lead, in my third lecture, to a novel account of the rational functional role that confers propositional conceptual content on declarative sentences.

II. Reasons and Practices

Discursive practices, in the sense I am focusing on, are those that enable practitioners to claim that things are thus-and-so. These are acts of *claiming*, both asserting and denying, which address specifically *propositional* contents. I began by considering a distinctive kind of contrast between act and content ('ing'/'ed'): the distinction between a *saying*, a speech act or discursive performance, and what is *said*, a content or meaning. Specifying that the sense of 'saying' that is in play here is *claiming*, *stating*, *saying or denying* ***that*** things are thus-and-so, is our first step in outlining a pragmatic theory. Linguistic expressions that can be uttered with that kind of *force* (asserting or denying, expressing attitudes of accepting or rejecting) and that kind of *content* (propositional) are thereby functioning as *declarative sentences*. That is why it is declarative sentences that are both uttered in speech acts of asserting or denying, and used to form the 'that'-clauses specifying the propositional contents asserted or denied.

So, we can functionally define declarative sentences as linguistic expressions whose free-standing utterance has the basic, default pragmatic force or significance of claimings, paradigmatically assertions. Is there a corresponding functional definition of that assertional force or pragmatic significance? The orienting idea I will pursue is that there is an internal conceptual relationship between *claiming* and *reasoning*. Asserting and denying are activities that take place in what Wilfrid Sellars calls "the space of reasons: of justifying and being able to justify what one says."[3] Discursive practices, in the sense of practices that accord some performances the practical significance of assertions or denials, are essentially (and not just accidentally) practices of giving and asking for reasons.[4] Making claims is part of a package of activities that includes rationally *challenging* claims by making further claims that contest them by offering reasons *against* them, and rationally *defending* claims by making further claims that

[3] *Empiricism and the Philosophy of Mind* §36.

[4] "What is it that we are *doing* when we assert, claim, or declare something? The general answer is that we are undertaking a certain kind of commitment... The idea is that assertings (performances that are overt undertakings of assertional commitments) are in the fundamental case what reasons are asked for, and what giving a reason always consists in. The kind of commitment that a claim of the assertional sort is an expression of is something that can stand in need of (and so be liable to the demand for) a reason; and it is something that can be offered as a reason... The idea exploited here, then, is that assertions are fundamentally fodder for inferences." (Brandom, *MIE* 1994: 167-168)

justify them by offering reasons *for* them. Of course not every claiming need be challenged or defended, but that claims are liable to challenges that bring with them justificatory obligations is a core feature of the genus.

Sellars offers the following historical metaconceptual context for this line of thought:

> Kant was on the right track when he insisted that just as concepts are essentially (and not accidentally) items which can occur in judgments, so judgments (and, therefore, indirectly concepts) are essentially (and not accidentally) items which can occur in reasonings or arguments.[5]

It is essential to assertings (including covert acts of judging) that they can both serve as and stand in need of reasons, can play the role both of premise and of conclusion in inferences. In fact Sellars draws (and attributes to Kant) semantic consequences from this holist, functionalist pragmatic order of explication. Thinking of saying that things are thus-and-so as 'descriptive' uses of expressions, he says:

> It is only because the expressions in terms of which we describe objects…locate these objects in a space of implications, that they describe at all, rather than merely label.[6] [CDCM §108]

What is disparaged as 'mere labeling' is a matter of an expression having determinate circumstances of appropriate application, so its proper production reliably differentially classifies what is being responded to. The surplus needed for genuine *conceptual* contentfulness, whose paradigm is *propositional* contentfulness, Sellars is claiming, is situation in a space of implications that determine also the appropriate *consequences* of application of the expression: what follows from the applicability of the concept. The suggestion is that propositions must be individuated at least as finely as the roles they play in a 'space of implications.'[7]

---

[5] "Inference and Meaning" [I-4], in Kevin Scharp and Robert Brandom (eds.) *In the Space of Reasons*: *Selected Essays of Wilfrid Sellars* [Harvard University Press, 2007].

[6] "Counterfactuals, Dispositions, and the Causal Modalities" §108.

[7] In fact, he thinks that semantically we must pay particular attention not only to which implications hold, but also to the *ranges of subjunctive robustness* of those implications. This is an important point, which I'll come back to in the third lecture, on the way to formalizing the Dummetian metaconcepts of circumstances and consequences of application.

I take these functionalist, specifically inferentialist, ideas to be right-headed and helpful. But in thinking about the relations between pragmatic claims about assertional *force* (what one is *doing* in saying that things are—describing things as—thus-and-so) and semantic claims about propositional conceptual *content*, it behooves us to be clear and careful about both the distinction and the relations between *inferential* moves or doings (justifying, challenging), on the pragmatic side of understanding assertional force in terms of practices of *reasoning*, on the one hand, and *implication* relations, on the semantic side of understanding propositional contents in terms of such relations, on the other hand. A famous argument due to Gilbert Harman highlights some crucial considerations. Harman provocatively claims that there is no such thing as deductive inference. If there were, he argues, surely a paradigmatic instance would be inferring from *p* and the conditional *if p then q* to the consequent *q*. But, he points out, thought of as a rule for *doing* something, in the sense of a general policy for belief-revision, that would be a bad one. For one might well have far better reasons or evidence *against q* than one had *for* either *p* or the conditional. And in that case one should surely *not* inferentially draw the conclusion *q* from one's commitment to *p* and *if p then q*. One should rather revise or reject one of those premises.

Harman's conclusion is that deductive logic determines implication relations, but that those *implication relations* do not determine, but only constrain proper *inferential practices*. We can put the point by saying that the fact that *p* and *if p then q* imply *q* means that it is inappropriate or incorrect to assert *p* and *if p then q* and to deny *q*. But if we find ourselves with good reasons for all those attitudes, the implication relations do not tell us what to *do* inferentially to repair that discordant constellation of doxastic commitments. The implicational relations among them do not settle what conclusion should be drawn, what assertional stance should be the outcome of an inferential move.

Taking on Harman's distinction between reason *relations* of implication or consequence between propositional contents and reasoning *practices* of justifying and challenging assertions and responding assertionally to such justifications and challenges puts us in a position to indicate the contours of a minimal model of discursive practice that can meet the dual basic criteria of adequacy of an account of the relations between pragmatic theory and semantic theory. These requirements are that it explain both how engaging in practices of making claims and reasoning

about them by inferentially justifying and challenging them can confer conceptual (paradigmatically propositional) content on assertible linguistic expressions, and how such contents can provide normative standards for assessments of the correctness of assertional and inferential performances. (These are what I earlier called the 'conferral' condition and the 'codification' condition on the relations between pragmatics and semantics.) The model is *minimal* in that all it aims to do is to characterize sufficient conditions on a practice for it to be a *discursive* practice—conditions, that is, that are necessary and sufficient for some performances to be accorded the practical significance of *claimings*. That requires understanding what is asserted by those acts of asserting or denied by those acts of denying (utterings of what thereby count as declarative sentences) to be propositional contents, which are the claimables claimed in claimings of the two kinds. Propositional contents are to be understood as what can both serve as and stand in need of reasons. Two principal criteria of adequacy are then that the minimal model of discursive practice make intelligible how engaging in that practice confers propositional contents, and how those propositional contents serve to normatively constrain the practice, in the sense of setting standards for assessment of the correctness of claimings, including those that play the role of justifying or challenging other claimings.

The idea for understanding the connections between implication relations and norms governing reasoning practices that I will pursue comes from the *bilateral* normative pragmatics of the consequence relations expressed by sequent-calculus turnstiles developed by Greg Restall and David Ripley. Put in my terms, they start with the idea of the *bipolarity* of doxastic attitudes. This is taking them to come in two essentially contrasting flavors: acceptance and rejection. We can think of them as co-ordinate with two flavors of the speech acts that express them: assertion and denial. Implication relations between a set of premises and a conclusion can then be understood in normative terms of assessments of the correctness or propriety of constellations of doxastic attitudes of the two kinds. In particular, for the premise-set of sentences Γ to imply conclusion sentence A (endorsing '$\Gamma|\sim A$') is for the overall position any speaker would be in if they *accepted* all of Γ and *rejected* A to be normatively "out of bounds," in Restall and Ripley's idiom. The two sides implicitly invoked in calling this definition 'bi*lateral*' are the two sides of the turnstile, which it proposes to treat differently with respect to the fundamental doxastic bipolarity of acceptance and rejection. For bilateralism understands the

significance for reason relations of what appears on the left-hand, *premise* side, in terms of doxastic *acceptance* and the significance for reason relations of what appears on the right-hand, *conclusion* side, in terms of doxastic *rejection*.

Restall and Ripley offer a definition of the central reason relation of implication that is couched in an avowedly *normative* pragmatic metalanguage. They are explicit about the normative character of the crucial concept of being 'out of bounds,' which entails incorrectness: a negative 'ought-to-be.' Our purposes will be served by offering a friendly amendment. For I take it that the notion of a 'position,' which is what is subject to normative assessment as in or out of bounds, is itself a normative notion. What is assessed is attitudes of acceptance or rejection (ultimately, of propositions), and those can usefully be thought of as the two basic kinds of doxastic *commitment*. A 'position' is then itself a constellation of normative statuses or attitudes, divided into commitments to *accept* and commitments to *reject*. For a set of premises Γ to *imply* a conclusion A is then for the set of *commitments* to accept everything in Γ and to reject A to be one to which one cannot be jointly *entitled*. 'Out-of-boundness' is a matter of *preclusion of entitlement* to all of a set of commitments.[8] Here it is worth noticing that the notion of doxastic commitment, whether to accept or reject, is an *atomistic* one. A speech act of assertion or denial can add a commitment to a position independently of what that position already contains. But entitlement is a *holistic* matter, assessed for whole positions, where entitlement to any one commitment depends on what else one is committed to.

In these terms, we can offer a rough characterization of the relation between rational implication *relations* and reasoning *practices*, that respects Harman's insight. Speech acts of claiming, like the doxastic commitments they express, must come in two flavors: assertions, expressing doxastic acceptance commitments, and denials, expressing doxastic rejection commitments. Basic reasoning practices govern two kinds of pragmatic significance some claimings can have relative to others: *challenging* a claim by offering reasons *against* it, and *justifying* a claim by offering reasons *for* it. The idea is that successfully challenging a claim

[8] In Chapter Five of *Articulating Reasons* I argue that any practice of giving and asking for reasons must distinguish something corresponding to the two normative statuses of commitment and entitlement, and that doing so provides a demonstrably more expressively powerful normative metalanguage than any single-sorted deontic metavocabulary, which work with one notion of appropriate/inappropriate or correct/incorrect.

voids the default entitlement doxastic commitments can be thought of as coming with, and successfully justifying a claim reinstates its status as a commitment entitlement to which is not precluded. How do reason relations, paradigmatically implication, figure in the doxastic practices of this model? The idea is that the pragmatic notion of the set of claimables Γ *being a reason for* the claim A appeals to the *implication* of A by Γ: the relation that holds when commitment to accept all of Γ precludes entitlement to deny A. Of course, as Harman would point out, whether offering Γ as a reason for A succeeds depends also, for instance, on entitlement to the premises in Γ.

What about reasons *against*? The *critical* dimension of rational practice—as opposed to the justificatory dimension—is articulated by another kind of reason relation: *incompatibility*. In the spirit of Restall's bilateral definition of the reason relation of implication, we can say that to give a reason *against* A is to make claims, Γ, acceptance of which precludes entitlement to *accept* A. There must be two basic kinds of reason relation because practices of reasoning require claims standing to one another both in the relation of being a reason *for* and in the relation of being a reason *against*.[9] In the most basic case, these are reasons to *accept* and reasons to *reject* (though there are also reasons for rejections, as well as against acceptances). We can take it that if Γ implies A, in the sense that commitment to *accept* all of Γ precludes entitlement to *reject* A, then Γ thereby *implicitly* provides reasons to *accept* A, since one has been precluded from being *entitled* to the only alternative commitment. Similarly, if Γ is incompatible with A, in the sense that commitment to accept all of Γ precludes entitlement to *accept* A, then Γ *implicitly* provides reasons to *reject* A, since one has been precluded from being *entitled* to the only alternative commitment.[10]

---

[9] Huw Price cogently assembles considerations showing that relations of doxastic incompatibility are essential to discursive practice, in "Why 'Not'?" [*Mind* V. 99, No. 394 (April 1990), pp. 221-238].

[10] MacFarlane ["In What Sense (If Any) Is Logic Normative for Thought" 2004] sharpens Harman's argument and question. MacFarlane's conclusion is:
"My own temptation is to go for a combination of Wo- and Wr+."
Wo- $=_{df.}$ you ought to see to it that if you believe A and you believe B, you do not disbelieve C.
Wr+ $=_{df.}$ you have reason to see to it that if you believe A and you believe B, you believe C.
Being precluded from entitlement to deny is implicit commitment to accept.
Inferring is explicitly acknowledging commitments that are implicit, in this specific sense of 'implicit'.
Our explicit notion of implication as reason for is of his type Wo-.
Our notion of implicit implication is of his type Wr+.

The paired reason relations of implication/ incompatibility accordingly articulate the pragmatic significance of being a reason for/against, which is an essential aspect of the fundamental bipolarity of acceptance/rejection. This is the same bipolarity that is expressed in a semantic metalanguage by the opposition true/false. These expressions of the basic bipolarity in the two idioms, pragmatic and semantic, are connected by the principle that acceptance is practically taking-true and rejection is practically taking-false. One of my principal aims in this lecture is to show that and how this platitude can be deepened substantially by focusing on the manifestation of the fundamental bipolarity at the level of reason relations among claimables, rather than of properties of those claimables—whether the *semantic* property of being true/false or the *pragmatic* property of being accepted/rejected by some interlocutor.

Reason relations of implication and incompatibility show up in this normative pragmatic model as systematic preclusions of entitlement to some constellations of doxastic commitments. So understood, reason relations give us a grip on the claimable propositional contents that stand in such relations. We can understand those semantic contents as the roles sentences expressing them play in structures of reason relations: both what they imply and what implies them, and what they are incompatible with. For these determine what is a reason *for* and *against* what, and so what speech acts of *challenging* and *defending* claims count as successful in altering which doxastic commitments various interlocutors are entitled to. In my third lecture, I will present a formal conceptual role semantics based on this idea of propositional contents as roles with respect to reason relations. For now it is enough to observe that the idea of propositional contents as standing in relations of implication and incompatibility as defined in the bilateral normative pragmatic vocabulary points to how the paired criteria of adequacy I called the *conferral* condition and the *codification* condition can be satisfied.

This minimal model of discursive practice shows how the use of declarative sentences by discursive practitioners can determine the reason relations that the claimables expressed by those sentences stand in to one another. The key question is which constellations of acceptances and rejections preclude entitlement to which others, according to the way practitioners keep track of entitlements in response to various claimings that they take or treat as reasons for or against

which others. Following the lead of Restall and Ripley's bilateral normative pragmatic definition of the reason relation of implication, we specify claiming and reason-giving practices in a richly normative metalanguage of commitments and preclusions to entitlement. Their example shows how we can explain what reason relations are, in those normative terms. By treating some claims as reasons for and against others, in the sense of what assertions and denials they count as challenges to and justifications or defenses of what others, practitioners institute reason relations among claimables, and so confer the kind of conceptual, propositional content articulated by the role claimables play in those reason relations.

The model of the role reason relations of implication and incompatibility play in practices of making claims and rationally challenging and defending them with reasons also explains how the reason relations we are thinking of as articulating propositional contents can be understood as providing normative standards for assessment of the correctness of linguistic performances. The pragmatic idea of rational challenges to claims as consisting in undertaking commitments that provide reasons against those claims, in the sense of being incompatible with them, and of rational defenses or justifications of claims as consisting in undertaking commitments that provide reasons for those claims, together show how what actually implies and is incompatible with what determines which challenges to or defenses of claims are successful in altering their entitlement status. The reason relations accordingly codify norms governing the making, challenging, and justifying of the propositionally contentful claims undertaken by assertions and denials.

Harman offered convincing reasons to distinguish rational *relations* of implication (and, by extension, incompatibility) and *practices* of giving and asking for reasons. I have suggested that a lightly tweaked version of Restall and Ripley's bilateral normative pragmatic definition of implication can be used to explain the relations between the two dimensions that Harman distinguishes. By explaining in normative pragmatic terms what reason relations of implication and incompatibility *are*, the bilateral account opens the way to an understanding of how *sayables* (in the sense of claimables) are related to *sayings* (in the sense of claimings). For we can appeal to the prospect of understanding those propositional claimables in terms of their role in reason relations.

The concept of reason relations, which the bilateral normative pragmatic account of implication has clarified, brings into view another concept, which in one way or another will be the topic of the rest of these lectures. That is the concept of a *vocabulary*. I will use that term in a technical sense, to refer to a certain kind of algebraic relational structure. By 'vocabulary' I mean an ordered pair <L,**I**>, whose first element is a domain of sentences, which we can call a *lexicon*, and whose second element is a specification of the reason relations holding between elements of the lexicon. More specifically, the reason relations are represented by a set of ordered pairs of sets of sentences from the lexicon. For a pair of sets of sentences to be included in the reason relations of a vocabulary means that the implication whose premises are the sentences included in the first element of the ordered pair and whose conclusions are the sentences included in the second element of the ordered pair is a *good* implication. (We can encode incompatibilities as well as implications in this format, by using Gentzen's convention that endorsing the goodness of an implication with an *empty* set of conclusions is to be read as taking the premise-set to be incoherent, in the sense that any element of it is incompatible with the rest.[11]) I take it that the (amended) bilateral definitions of implication and incompatibility tell us, in a deontically two-sorted (commitment/entitlement) normative pragmatic metalanguage, what it is, in broad outlines, for a linguistic community to *use* a vocabulary. It is for their practices of making, challenging, and defending claims expressed by sentences of the lexicon of the vocabulary to be normatively governed by the reason relations of that vocabulary. The vocabulary articulates the claimables—sentences playing roles with respect to reason relations—whose use in claimings the minimal model of discursive practice explicates.

What we have seen so far is how vocabularies look from the point of view of pragmatics: both how the vocabulary can be elaborated from what discursive practitioners do (according to the minimal model), as expressed in a relatively rich normative pragmatic metalanguage, and how the vocabulary can be understood as providing a normative standard for assessments of the correctness of discursive performances of claiming and challenging and defending claims. I turn

[11] In fact I am thinking of the presence of an ordered pair of sets of sentences <Γ,Δ> in the second element of the relational structure that is a vocabulary as indicating that the multisuccedent sequent Γ|~Δ is good in the bilateral sense that commitment to accept all of Γ precludes entitlement to deny all of Δ. But this detail does not matter for my story at this point.

next to consider how vocabularies, in this regimented sense of a lexicon plus a set of reason relations defined on that lexicon, look from a more orthodox *semantic* point of view.

III. Truthmaker Semantics

The most sophisticated and expressively powerful contemporary representational formal semantic framework is Kit Fine's truthmaker semantics. It begins with a metaphysical picture of what there is to be represented semantically. That universe consists of a structured collection of what he calls 'states.' The formal apparatus is as noncommittal as possible about what these consist in, but states are meant to include such ways things could be as Pittsburgh's being to the West of New York City and snow being white. The universe of states is thought of as having two sorts of structure: mereological and modal. On the mereological side, some states are to be understood as being *parts* of others. More formally, there is a *fusion* operation that maps any set of states into a whole comprising them. This defines the part-whole relation: state A is part of state B just in case B is the result of fusing A with some other states. On the modal side, the universe of states is partitioned into *possible* and *impossible* states.

Mereologically and modally structured state spaces generalize the metaphysics of possible worlds in a number of important ways. Possible worlds show up in this framework as maximal possible states: possible states such that every other state is either a part of that state or incompatible with it, in the sense that fusing it with the world-state yields an *im*possible state. Situation semantics had already shown the expressive advantages of building such wholes out of smaller parts, rather than getting the partial ones by analyzing whole worlds. On the modal side, state spaces in general include multiple *im*possible states, where the possible worlds setting in effect has only one. On the mereological side, various structural conditions can be put on the fusion operation, for instance, requiring that all the states that contain any impossible state are themselves impossible—that is, that the result of fusing any state with an impossible state is always an impossible state. Like the existence of multiple impossible states, the capacity to consider different kinds of mereological structures is a major degree of freedom in the apparatus, enhancing the expressive power of the truth-maker framework.

This metaphysical specification of what is there to be represented is then married to a flexible and powerful semantics. An interpretation function assigns each declarative sentence to a pair of

sets of states, thought of as the (exact) *truth*-makers and *falsity*-makers of that sentence. Rather than simply defining one of these sets in terms of the other, one can put various explicit structural constraints on the sets of verifiers and falsifiers that are assigned to declarative sentences as their semantic interpretants. One might be tempted to require that they be disjoint: no possible state is both a truth-maker and a falsity-maker of any sentence. Fine requires rather that the fusion of any truth-maker with any false-maker of the same sentence must be an impossible state. He calls this condition *Exclusivity*. It entails the cognate, but usefully different, requirement that any states that are both truth-makers and false-makers of the same sentence be impossible states. Some statements, say "All cows are made of glass," and "This photon has a mass of 500 kilograms," might have only impossible truth-makers—but they are not required to have the *same* impossible states as truth-makers. The combination of the mereological and modal fineness of grain of the underlying metaphysics and keeping separate books on the truth-makers and falsity-makers that semantically interpret sentences results in a hyperintensional theory of meaning, which makes many more distinctions than its possible-worlds predecessor.

Together, Fine's modal-mereological metaphysics and truthmaker semantics underwrite a striking realism about the propositional contents expressed by declarative sentences. Such contents are just pairs of sets of states that meet whatever structural conditions we impose on such pairs to make them eligible to serve as truth-makers and false-makers of sentences—paradigmatically, Exclusivity, which requires that all fusions of elements of the first set with elements of the second set be impossible states. A proposition, for Fine, is any pair of sets of states meeting that condition, since it is eligible to serve as the interpretant of a sentence. Even in the metaphysically implausible case where there is only a countably infinite number of states, there will be uncountably many pairs of sets of them meeting the minimal structural condition for propositionality—so, far more than any natural or formal language in the ordinary sense can have sentences to express.

Further, those worldly propositions, understood as set-theoretic and mereological constellations of possible and impossible states, stand to one another in relations of consequence and incompatibility. Fine offers two principal ways one might define consequence and counts it a virtue of the system that there are such alternatives. He says that a set of sentences Γ *entails* a

conclusion A in case every verifier of all the premises in Γ is also a verifier of A. He says that A is a consequence of Γ in the sense of *containment* if and only if every verifier of A includes as a part a verifier of all of Γ and every verifier of all of Γ is a part of a verifier of A. Corresponding definitions of incompatibility are not far to seek. Here I want to offer a friendly amendment. We might notice, to begin with, that Fine's definitions of consequence relations do not make anything like full use of the mereological and modal innovations that principally distinguish his framework from the possible worlds semantics it develops from and improves upon. Taking it that premise-set Γ entails conclusion A just in case all the verifiers of all of Γ are verifiers of A just translates the set-theoretic inclusion criterion of consequence from the possible-worlds setting, without adding anything of substance to it. His notion of containment exploits the mereological structure of his metaphysics, but not its modal structure.

My collaborator and coauthor, Ulf Hlobil shows us how to do better.[12] We can take our cue from Fine's Exclusivity condition relating verifiers and falsifiers of the same sentence (so of the same proposition). It requires that every fusion of any verifiers and any falsifiers be an impossible state. Hlobil suggests that we take Γ to *imply* A just in case every fusion of any verifiers of all of Γ with any falsifier of A is an impossible state. Exclusivity of verifiers and falsifiers is then manifested as the Reflexivity of the consequence relation: the principle that every premise implies itself. Like Exclusivity, this definition of a notion of consequence appeals both to the mereological and to the modal structure of the universe of states from which the semantic interpretants of sentences are drawn. (The corresponding notion of incompatibility requires that the fusion of any verifiers of all of Γ with any *verifier* of A be an impossible state.)

The key point is that this semantic definition of implication lines up perfectly with the bilateral *pragmatic* definition of implication.
Hlobil's version of consequence (implication) in truth-maker semantics is:

1. Γ implies A iff any fusion of a state that verifies all the members of Γ with a state that falsifies A is an impossible state.

[12] Hlobil, U. (2022a). The laws of thought and the laws of truth as two sides of one coin. *Journal of Philosophical Logic*, 52:313–343.

The Restall-Ripley normative pragmatic reading of implication is:

2. Γ implies A iff any position that includes accepting all of Γ and rejecting A is normatively incoherent or "out of bounds"—as we have read it: one cannot be entitled to such a constellation of commitments.

And similarly for incompatibility:

3. Γ is incompatible with A ⇔ the state resulting from *fusion* of any *verifiers* of all the members of Γ with any *verifier* of A is an *impossible* state,
4. Γ is incompatible with A ⇔ the position resulting from concomitant commitment to *accept* all of Γ and to *accept* A is normatively *incoherent* ("out of bounds")—a constellation of commitments to which one *cannot* be entitled (entitlement to which is precluded).

Indeed, Hlobil proves that with these definitions, the reason relations defined semantically in Fine's truth-maker setting are *isomorphic* with those defined pragmatically in the bilateral normative setting.

IV. Representation and Reason Relations

This isomorphism is the result I have been building up to in this lecture so far. It points to a structure that is common to or amphibious between a *pragmatic* account of what one is *doing* in saying or claiming that things are thus-and-so, on the one hand, and a *semantic* account of the sayable or claimable propositions that are *said* or *claimed*. For both on the pragmatic side of speech acts and on the semantic side of conceptual contents the isomorphism reveals *reason relations*, bivalently distinguished into relations of implication and incompatibility. The isomorphism shows that the *very same* reason relations can be understood as both defining reasons *for* and reasons *against* suitable for underwriting practical *defenses of* and *challenges to* doxastic commitments, and also as defining relations among worldly propositions. From the pragmatic point of view, both kinds of reason relation appear as the incompatibility of a set of doxastic commitments (to accept or reject), in the sense that no-one can be normatively entitled to all of those commitments. From the semantic point of view, both kinds of reason relation appear as the incompatibility of a set of propositions, in the sense of the modal noncompossibility of the worldly states those propositions comprise.

The common topic revealed earns the right to be called "reason" relations in virtue of its role in a pragmatic account of reasoning practices. These include not only rationally challenging and defending commitments by giving reasons for and against them by making further claims that stand to them in relations of incompatibility and implication, but also inferring in the sense of explicitly acknowledging an implicit commitment by accepting what one is precluded from entitlement to deny, or denying what one is precluded from entitlement to accept. Hlobil's isomorphism articulates a precise sense in which those *very same* abstract reason relations can be discerned in the relations among worldly propositions, understood according to Fine's semantics and the modal-mereological metaphysics of objective reality on which the truthmaker semantics is based. The reason relations so discerned are 'abstract' in the technical sense: they are the result of treating an equivalence relation as an identity. (Frege's model in the *Grundlagen* is abstracting directions of lines from the equivalence relation of being parallel.)

The pragmatic and semantic theories I have sketched offer substantive specifications and accounts of relations of implication and incompatibility. But they are very different. Our version

of the bilateral pragmatic account of the use of sentences is in terms of normative statuses of *commitment* and *entitlement*, and practical doxastic attitudes of *acceptance* and *rejection*. It tells us what it is for a discursive community practically to take or treat various implication and incompatibility relations as holding, by implicitly acknowledging them as norms governing their practices of making claims and challenging and defending them with reasons. This pragmatic account of the subjective appearance of reason relations is mirrored by a semantic account of the objective metaphysical realities, couched in terms of the alethic impossibility of certain mereologically fused states, that determine what propositions in fact follow from and are incompatible with which others. The isomorphism proves that the two metalinguistic accounts can offer perspectives on a single common more abstract structure of implication and incompatibility.

So the isomorphism between these two ways of specifying reason relations precisely determines their common topic: the relations of implication and incompatibility that were informally introduced by Harman's argument. Astonishingly, it allows us to see that and how those reason relations articulate a structure common to assertional force and propositional content. I introduced the term 'vocabulary' to refer to relational structures consisting of a domain of sentences, the lexicon, and a set of reason relations on that domain. The formal representation of reason relations I will use in specifying vocabularies is a set of pairs of sets of sentences, interpreted as the good implications relating premise-sets of sentences to conclusion-sets of sentences.[13] Then we can see both the bilateral normative pragmatic theory and the truthmaker modal-mereological semantic theory as offering explanatory accounts of reason relations, and hence of vocabularies in this sense—even as we are pointed to a broader model in which the things that stand in reason relations are thought of not as sentences, but as propositions in Fine's sense of pairs of sets of objective states satisfying Exclusivity. The next two lectures explore and exploit these newly precise concepts of reason relations and vocabulary. In

[13] In order to facilitate later revealing illuminating connections (both logical and semantic) between vocabularies and sequent calculi, we read the premise sets conjunctively and the conclusion sets disjunctively—that is, Gentzen-wise rather than Tarski-wise. But that difference does not make a difference at this point in my story. As previously noted, incompatibility relations are encoded by empty conclusion-sets (which, in the presence of structural rule of weakening on the right, on the disjunctive multisuccedent reading becomes equivalent to having the whole lexicon as a conclusion).

particular, I will investigate the notion of conceptual content as consisting in the functional role some bearer (sentential or otherwise) plays with respect to the reason relations in a vocabulary.

Already we can see how focusing on reason relations illuminates what pragmatic and semantic theories owe each other, the criteria of adequacy they set for each other. The Hlobil isomorphism between (suitably tweaked versions of) Fine's truth-maker representational semantics and (our deontically two-sorted version of) Restall and Ripley's bilateral normative pragmatics supplies an answer to a question Fine's framework by itself does not. For it begins to tell us what practitioners must *do*, how they must *use* expressions, in order to confer on them the conceptual contents Fine assigns them in terms of truth-makers and falsity-makers, up to isomorphism of reason relations. To associate verifiers and falsifiers with expressions as their semantic interpretants, practitioners must use those expressions according to the bilateral pragmatics, distinguishing in practice between constellations of commitments to accept and reject claimables that are normatively "in bounds" and those that are normatively "out of bounds." That includes expecting anyone who is precluded from being jointly entitled to the doxastic commitments they have undertaken practically to acknowledge the obligation to alter those commitments so as to repair the situation and find their way back in bounds.

This account does not explain what it is to use a sentence so as to confer on it a relation to one pair of sets of states rather than another, if those pairs of verifiers and falsifiers are incompatible with, or imply and are implied by, the same propositions. The isomorphism is only up to reason relations. Fine would still owe an account of how the more fine-grained semantic relations to states that his semantic interpretation functions appeal to can be established by the use of linguistic expressions. Still, being able to say exactly how assertional uses of sentences implicitly acknowledge the normative significance of reason relations, which we can also understand in his semantic-cum-metaphysical terms, represents real progress on this front. For it at least shows how using sentences assertionally is intelligible as treating those sentences as expressing propositional meanings, insofar as conceptual contents are identified with roles in implication and incompatibility relations.

It is important that the isomorphism, and so the correspondence between representings and representeds, is specified to begin with not at the level of linguistic *sentences* and worldly *facts*, but at the higher level of reason relations. This is the very top of the hierarchy Sellars describes Kant as having arrived at by inverting the traditional bottom-up order of explication. That is, on the linguistic side it is at the level of *meaning*, not of *truth*. The common structure we have

discerned does not depend on what anyone is actually doxastically committed to, on the pragmatic side of representings, nor on what states are actual, on the semantic side of represented reality. It is not a correspondence theory of truth. Rather, the sort of conceptual realism it underwrites is a transcendental presupposition of the possibility of correspondence theories of truth.[14] One reason the possibility of understanding representation at the level of reason relations has not been sufficiently explored is that people have not been working with a sufficiently developed pragmatics. Another is insufficient appreciation of the lessons taught by the history of the concept of representation.

[14] The idea of "coherence theories of truth" was always the result of misunderstandings of holistic theories of *meaning*. Truth of sentences as correspondence to reality is a local property, appropriate to atomistic categories of resemblance rather than the holistic categories of representation presupposed at the level of meaning, on which it turns out to depend. At the level of reason relations, we see new possibilities emerge for combining elements of coherence and correspondence in a semantic theory.

V. From Atomistic Resemblance to Holistic Representation: Bimodal Conceptual Realism

For this reason, it is worth standing back a bit and looking at how the isomorphism between pragmatics and semantics at the level of reason relations illuminates the specifically *representational* semantic dimension of discursive practice. The concept of representation is a distinctively modern one, essentially due to Descartes. The classical philosophical tradition understood the relation between appearance and reality (in that sense, mind and world) on the model of *resemblance*, whose paradigm is the relation between a picture and what is pictured. The idea is that a picture (or idea) resembles its object, and so is veridical, insofar as it *shares* visual *properties* of shape or color with it, and is a misleading appearance insofar as does not in this sense resemble what it pictures. Descartes saw that this model begins to break down when it is applied to the theories of the new science of his time. Copernicus claimed that the reality behind the appearance of a stationary Earth and Sun revolving around it is a rotating Earth and stationary Sun. Any resemblances there run the wrong way. And it gets worse. The appearances Galileo found most veridical and (so) useful for reasoning about physical reality have periods of time appearing as the lengths of lines, and accelerations as the areas of triangles. What properties are shared there to underwrite a resemblance? And in the case of Descartes's own analytic geometry, the relation between the equations $x^2+y^2=1$ and $x+y=1$ on the one hand, and the circle and line they determine, on the other, is certainly not one of resemblance in the traditional sense.

Descartes sees that a more abstract concept of representation is required to handle these cases, since the more intuitive notion of resemblance ceases to be useful just when it is most urgently needed: in explaining the sense in which the new science offers *better*, more veridical appearances of physical reality than common sense. He did not go on to offer a useful *account* of this more general and abstract relation of representation, however. He took it to be essentially a brute fact that the world contains two kinds of things, representings, which are by nature *tanquam rem*, of or about things, and extended things, which are merely represented or representable.

Spinoza (whose first book was on Descartes) figured out the most basic features of the concept of representation that was implicit in the motivating paradigm of analytic geometry. The key is that, as he puts it, "**the order and connection of ideas is the same as the order and connection of things**."[15] Equations can represent geometrical figures because the whole *system* of equations is isomorphic to the whole *system* of figures—with, for instance, simultaneous solutions of equations corresponding to intersections of lines. That is why algebraically manipulating equations is intelligible as reasoning about geometrical figures. Given the global isomorphism—the "order and connection" of linear strings of symbols that is the same as the "order and connection" of extended plane figures—the equation '$x^2+y^2=1$' can play the same functional role in the world of equations that the circle it thereby counts as representing plays in the world of geometrical figures, with simultaneous solutions of equations algebraically playing the role of geometrical intersection of lines.

According to this story, the resemblance model was not wrong to take the sharing of properties to be essential to the *of*-ness invoked by talk of appearances *of* material reality. Its mistake, the source of its expressive limitations, was to restrict attention to *local* properties, conceived *atomistically*: properties elements of picturings and of what is pictured could have regardless of what properties *other*, systematically related elements had. Spinoza saw that the wider scope of the new representational model is due to the *holistic* character of its appeal to *global* isomorphisms, which make visible *functional* correlations between items in the two systems that might have quite different atomistic material properties. The new, more abstract and expressively powerful *representational* model of the intentional nexus between appearance and reality develops the older, more concrete *resemblance* model by shifting attention to the larger relational structures whose individual elements can be understood to play the functional roles of representing and represented in virtue of the global isomorphism of those structures. Representings and representeds are still understood to share properties—but properties of a new, functional kind, intelligible only globally and holistically, that is, *relationally*: in terms of relations to other representings (or representeds, respectively).

[15] *Ethics* II, Prop 7.

Spinoza elaborated this functionalist, holist conception of the representational relations between mind and world in a rationalist spirit. For him the abstract, systematic "order and connection" that is common to the two poles of the intentional nexus consists of *rational* relations—just as our isomorphism is at the level of reason relations. With the recollective wisdom of hindsight, we can see that Spinoza was hobbled in developing his rationalist, holist insight by two connected mistakes, both of which Kant would rectify. First, he still followed the Aristotelian tradition of thinking about the items that most proximally stand in representational relations to one another as *particulars*, thoughts and things (finite modes of substances), rather than as subjective judgings and objective judgeables, that is, *propositions*. Second, he failed to appreciate the distinctive *normative* character of the "order and connection of ideas." This is the Kantian thought that applying concepts or ideas in judging is undertaking a kind of *commitment*, something the judger is *responsible* for, something that can stand in entitlement relations of licensing and prohibition to other such commitments. Spinoza thought of the relations making up the "order and connection" common to thoughts and things in terms of *'necessity.'* He understood necessity in alethic modal terms, and found its paradigm in the lawful regularities natural science was beginning to codify. By contrast, Kant used the term 'necessary' ('notwendig') to mean "governed by a rule." In addition to government by rules in the alethic modal sense of laws of nature, he acknowledged *practical* necessities where deontic normative rules become *normatively* binding only when endorsed by autonomous agents.

I mention these mighty dead philosophers both to provide context for and to emphasize some radical features of the conception I have been articulating. The fundamental structural identity between features of the use of linguistic expressions and their objective correlates is at the level of reason relations, of implications and incompatibilities, not of sentences and facts nor terms/predicates and particulars/relations. Declarative sentences and worldly propositions are correlated, just insofar as they play the same role with respect to reason relations. The roles with respect to reason relations shared by the sentences used to make assertions and the worldly propositions they express are conceptually articulated by the reason relations they stand in to others of their kind. This is a kind of *conceptual realism*, in that conceptual contents are to be found on both sides of the intentional nexus, in the world as well as in the practices of linguistic communities. This distinctive sort of realism is made possible by a *non-psychological*

conception of the conceptual as consisting in role with respect to reason relations. According to this conceptual realism, worldly propositions, as pairs of sets of states that meet the modal Exclusivity condition, can themselves be bearers of conceptual content, by standing to one another in relations of implication and incompatibility. In this sense, the objective world would have been conceptually structured by relations of consequence and incompatibility even if there never had been talkers-and-thinkers to reason according to those relations.

The roles with respect to reason relations that are shareable between items caught up in discursive practices of claiming and defending and challenging claims, on the one hand, and constellations of worldly states, on the other hand, can be thought of as *rational forms*, in a recognizably neo-Aristotelian sense. They are *rational* forms precisely in being roles things play in structures of reason relations. They are essentially *modal* forms. For both essentially appeal to a notion of preclusion: the **im*propriety*** ("out of boundness") of a collection of concomitant commitments, or the **im*possibility*** of a state resulting from the fusion of other states. In both cases, consequence is a matter of a kind of necessitation, and incompatibility of a kind of exclusion. The isomorphism shows that the modal relations can correspond exactly.[16] But the *kinds* of modality involved in the pragmatics of representing and in the metaphysics of the representeds in the semantics are systematically different.

The modality that articulates the reason relations implicit in the use of declarative sentences is a *deontic* modality, while that articulating the reason relations implicit in the modalized mereological universe of states is an *alethic* modality. On the pragmatic side of claimings, the claim that the coin is made of copper is materially incompatible with the claim that the coin is an electrical insulator. The modal ruling-out involved in this kind of incompatibility is *normative*:

[16] Hilary Putnam argues (beginning in "Models and Reality" *Journal of Symbolic Logic* vol. 45, No. 3 (Sep. 1980), pp. 464-482) that isomorphisms come cheap. For instance, given any two physical systems, there is some vocabulary in which they will have a common structural description. This objection does not show that the isomorphism between reason relations as specified in a bilateral deontic normative pragmatic metavocabulary and reason relations as specified in a truthmaker alethic modal semantic metavocabulary is trivial. First, we start with two vocabularies for specifying reason relations that are fixed in advance and independently motivatable as characterizing the use and meaning of linguistic expressions, respectively. Second, our isomorphism is *modally robust*—modulo the difference between deontic and alethic modalities. As will be discussed in my third lecture, our isomorphism extends to the *ranges of subjunctive robustness* of the implications and incompatibilities specified. The resulting isomorphism is not vulnerable to the sort of extensional, model-theoretic trivializing counterexamples Putnam appeals to.

one cannot be *entitled* to commitments to accept both these claimables. It is *possible* to do so, but not *appropriate* or *correct*. The state consisting of the coin's being made of copper and the state consisting of the coin's being an electrical insulator are incompatible in the different, alethic modal sense that the combination of them that is their metaphysical fusion is *impossible*. Both the order and connection of ideas and the order and connection of things consist of *modally robust* reason relations: in the one case deontic normative, and in the other case alethic modal. The view I am recommending is accordingly a *bimodal* conceptual realism.

Although the early Modern philosophical tradition I have sketched culminates in a top-down order of explication, it never fully came to terms with the *holism* about conceptual contents implicit in it. It did not because of the impoverished conception of *relations* bequeathed it by Aristotelian and Scholastic logic. Substances and their individual modes are foundational for Descartes and Spinoza. Leibniz denies the reality of relations entirely, treating them as "well-founded phenomena" emergent from the co-ordination of monadic perspectives. Even Kant treats relations as transcendentally ideal, understanding all relations as the products of synthetic activities of the transcendental subject operating on the diverse manifold of particular representings delivered by the senses. Hegel was the first to pursue the radical top-down, holistic program of understanding both the ideas and things that stand in the relations that articulate the "order and connection of ideas" and the "order and connection of things" *relationally*: as *constituted by* those reason relations. (He took his Romantic predecessors' dominant metaphor of *organic* unity as anticipating this insight, while misunderstanding its essentially conceptual character.) The question of whether such a holistic ontological conception is ultimately so much as intelligible has been with us ever since. In my third lecture, I will discuss how a top-down, relations-first order of understanding of propositional conceptual contents in terms of reason relations can be worked out in detail.

VI. Conclusion

My principal concern in this lecture has been to introduce the concept of reason relations of implication and incompatibility, along with the closely related concept of a vocabulary, which is a lexicon of sentences together with a set of reason relations among them. To do that, I pursued a stereoscopic triangulation strategy: comparing and contrasting how implication and incompatibility can be characterized in a bilateral pragmatic metavocabulary expressing what interlocutors are *doing* in using declarative sentences to make claims and challenge and defend them, with how they show up in a truthmaker semantic metavocabulary expressing what interlocutors are *saying*, the propositions they are expressing by engaging in discursive practice. Appealing to a result due to Ulf Hlobil, I showed how to define reason relations *isomorphically* in the bilateral pragmatics and the truthmaker semantics. The isomorphism both secures a common (abstract) topic and underwrites a recognizably *representational* relation at the level of reason relations, connecting the norms governing challenging and defending doxastic commitments to modally robust relations of consequence and incompatibility among worldly propositions.

My next lecture addresses two related topics: the structure of reason relations in general and the relation of distinctively *logical* relations of consequence and inconsistency to reason relations in general. The first question is: What are the minimal structural conditions on reason relations that are compatible with the concordance between use and meaning given definite content by the fundamental isomorphism of pragmatically and semantically specified reason relations? We can approach the issue of the ultimate structure of reasons by investigating how robust the isomorphism we are using to pick out reason relations is under variations in their structure. The second question is: What are the relations between *reasons* and *logic*? Here I will offer an account of logical vocabulary as distinguished by its characteristic *expressive* role of making reason relations explicit. This expressivist account sets criteria of adequacy for an ideal logical vocabulary: that it have the expressive power to make explicit any and all reason relations, whatever their structure. Again reporting results from our book, I will present a logic that can be shown to maximally satisfy this expressive ideal. The third lecture completes the investigation of reason relations by presenting a pure model-theoretic semantics of the propositional

conceptual roles expressed by sentences as defined solely by reason relations, independently of whether they are specified in pragmatic deontic normative or objective alethic modal terms. I will then survey the virtues of that implication-space semantics for arbitrary vocabularies—including massively substructural ones—beginning with a sound and complete semantics for the logic introduced in the previous lecture. My overall aim is to use the very spare and simple representation of reason relations in vocabularies to illuminate discursiveness from a number of perspectives: to begin with, pragmatics and representational semantics, and then from the points of view of logic and an implication-space semantics of pure conceptual roles.

End of Lecture 1

Abstract of Lecture 2:

Logic and the Structure of Reasons

What is the relation between material reason relations of implication and incompatibility in general and specifically *logical* relations of consequence and inconsistency? Unlike traditional logically valid reasons, good reasons more generally can be defeasible. In spite of this structural mismatch, it is possible to use the tools of logic to make explicit even the most radically substructural reason relations. The logic NMMS (for Nonmonotonic, Multisuccedent), introduced in a sequent-calculus metalanguage, is expressively complete for the reason relations of (almost) any base vocabulary of material reason relations among logically atomic sentences. Perhaps surprisingly, it is just a version of classical logic, and the purely logical consequence relation it determines is monotonic and transitive, even though the reason relations it codifies and expresses explicitly are not.

Pufendorf Lectures 2025: **Vocabularies of Reason**

Handout for Lecture II

Logic and the Structure of Reasons

Bob Brandom

I. The Structure of Reason Relations

What is the relation between *material* reason relations of implication and incompatibility and *formal logical* relations of consequence and inconsistency?
The former includes implications such as
Pittsburgh is to the West of New York |~ New York is to the East of Pittsburgh.
It is raining |~ The streets will be wet.
These are good because of the content of the nonlogical concepts involved.

Material reason relations facially do not have the same structure as formal logical ones.
Logical consequence and inconsistency are monotonic and transitive.
Tarski's view is that the essential structure of consequence relations is that they are *topological closure operators*. In our terms, he takes a vocabulary to be a lexicon L of sentences, and a two-place relation between sets of sentences pairing each subset $X \subseteq L$ with its consequence-set Con(X), satisfying these conditions:
Containment (CO): $X \subseteq Con(X)$.
Monotonicity (MO): $X \subseteq Y \Rightarrow Con(X) \subseteq Con(Y)$.
Idempotence (CT): $Con(Con(X)) = Con(X)$.
(These are variants of the Kuratowski axioms for topological closure operators.)

But materially good implications can be nonmonotonic:
Tweety is a bird |~ Tweety flies.
Tweety is a bird and Tweety is a penguin # Tweety flies.
And we can turn failures of monotonicity into failures of transitivity:
Tweety is a penguin |~ Tweety is a bird, Tweety is a bird |~ Tweety flies
These two premises do *not* entail: Tweety is a penguin |~ Tweety flies.

Cautious Monotonicity (CM): $\dfrac{\Gamma \mid\sim A \quad \Gamma \mid\sim B}{\Gamma, B \mid\sim A.}$

Cumulative Transitivity (CT): $\dfrac{\Gamma, B \mid\sim A \quad \Gamma \mid\sim B}{\Gamma \mid\sim A.}$

If Γ|~A and G∈Γ, then G is part of the *explicit* content of Γ, and A is part of the *implicit* content of Γ—in the literal sense of being *implied by* Γ.
Explicitation is moving a sentence from the conclusion side of an implication to the premise side. This is the paradigmatically rational activity of explicitly acknowledging the consequences of one's commitments.
Since CM says that explicitation never *subtracts* consequences from a premise-set, and CT says that explicitation never *adds* consequences to it, together they entail that ***explicitation* is *inconsequential***. But explicitation is *not* always inconsequential. Example: database of observations plus theory as inference engine to extract consequences.
So some implication relations are *hyper*nonomonotonic: neither topologically closed nor explicitationally closed.
In these radically substructural settings, extracting consequences by explicitation is path dependent: it matters in what order one extracts consequences. There need be no such thing as *the* unique set of consequences of a premise-set.

II. Expressivism and Logic

The *reasons question* about logic asks: What is the relation between logic and reasons?
Logicism about reasons: Good reasons are always *logically* good reasons.
Structural logicism about reasons:
The structure of reason relations is the topologically closed structure of *logical* reason relations.
Logical expressivism: The expressive role characteristic of logical vocabulary is to make explicit reason relations of implication and incompatibility.
Expressive goal of *structural universality*: The reason relations codifiable by expressively ideal logical vocabulary include not only topologically open (nonmonotonic and nontransitive) ones, but also explicitationally open ones (hypernonmonotonic reason relations, in which even CM fails).
The expressively ideal logic is *universally LX*: it can be conservatively *elaborated from* (L) and is *explicative of* (X) any constellation of reason relations whatsoever, regardless of its structure.

Sequent calculi consist of metainferential rules that show how to extend the reason relations to correspond to an extension of vocabulary by introducing logical operators.
The operational rules do that, in the context of the structural rules.

Deduction-Detachment (DD):

$$\frac{\Gamma, A \mid\sim B, \Delta}{\Gamma \mid\sim A \rightarrow B, \Delta.}$$

Bidirectional Meta-Inference Line

Incoherence-Incompatibility (II):

$$\frac{\Gamma, A \mid\sim \Delta}{\Gamma \mid\sim \neg A, \Delta.}$$

Bidirectional Meta-Inference Line

There are many different formulations (for instance, in sequent calculus MVs) that are equivalent in that they all specify the consequence relation of classical logic. Under substructural conditions, they come apart, and specify different consequence relations.
One of them—which is just a formulation of classical logic, under strong structural conditions—degrades gracefully when we go substructural.
It remains universally LX, and expressively complete in a very strong sense.

*N*on*M*onotonic,*M*ulti-*S*uccedent logic, or NMMS for short, has three remarkable properties.

1. It is *expressively complete* in an unprecedentedly strong sense. We show how to associate each sequent or set of sequents whose premise-set and conclusion-set consist of logically complex sentences with a set of sequents in the base vocabulary, which relate only the logically atomic sentences that occur in them, such that those sequents from the logical supervocabulary hold in all and only the NMMS-elaborations of bases in which just those atomic sequents hold. In this clear sense, the logically complex sequents *say that* the corresponding logically atomic sequents hold. It is LX for its base vocabularies.
2. The second remarkable property of NMMS is that it is fully tolerant of open-structured or radically substructural base vocabularies. This feature has to do with the *elaboration* dimension of LX-ness rather than the *explication* dimension, as the first point did. NMMS conservatively extends logically atomic base vocabularies that are nonmonotonic and nontransitive, those in which Cautious Monotonicity fails, and even those for which Containment fails. It is *universally* LX.
3. The third remarkable property of NMMS is that it is essentially just classical logic. In the fully topologically closed setting defined by Gentzen's full set of structural rules, NMMS yields exactly the same logically valid sequents as Gentzen's sequent-calculus version of classical logic, LK. NMMS is supraclassical when applied to any base vocabularies that include all instances Containment, and yields exactly the classically valid implications and incompatibilities if it is applied to base vocabularies *all* of whose implications are instances of Containment

Connective Rules of NMMS:

L¬: $\dfrac{\Gamma\mid\sim\Delta,A}{\Gamma,\neg A\mid\sim\Delta}$ R¬: $\dfrac{\Gamma,A\mid\sim\Delta}{\Gamma\mid\sim\Delta,\neg A}$

L→: $\dfrac{\Gamma\mid\sim\Delta,A \quad B,\Gamma\mid\Delta \quad B,\Gamma\mid\sim\Delta,A}{\Gamma,A\rightarrow B\mid\sim\Delta}$ R→: $\dfrac{\Gamma,A\mid\sim B,\Delta}{\Gamma\mid\sim A\rightarrow B,\Delta}$

L&: $\dfrac{\Gamma,A,B\mid\sim\Delta}{\Gamma,A\&B\mid\sim\Delta}$ R&: $\dfrac{\Gamma\mid\sim\Delta,A \quad \Gamma\mid\sim\Delta,B \quad \Gamma\mid\sim\Delta,A,B}{\Gamma\mid\sim\Delta,A\&B}$

L∨: $\dfrac{\Gamma,A\mid\sim\Delta \quad \Gamma,B\mid\sim\Delta \quad \Gamma,A,B\mid\sim\Delta}{\Gamma,A\vee B\mid\sim\Delta}$ R∨: $\dfrac{\Gamma\mid\sim\Delta,A,B}{\Gamma\mid\sim\Delta,A\vee B}$

Material is from Chapters 2 and 3 of Ulf Hlobil and Robert Brandom
Reasons for Logic, Logic for Reasons: Pragmatics, Semantics, and Conceptual Roles [Routledge, 2024].

Vocabularies of Reason
2025 Pufendorf Lectures
Lund University

Lecture II

Logic and the Structure of Reasons

I. Three Versions of a Basic Discursive Bipolarity

In my first lecture I introduced the concept of reason relations of implication and incompatibility, motivated by Harman's argument for the need to distinguish such rational *relations* from inferential *practices*. I showed that such consequence and exclusion relations can be understood as amphibious between *pragmatic* accounts of what one is *doing* in *using* declarative sentences to say something in the sense of asserting or denying it, and *semantic* accounts of what one *saying*, the proposition being accepted or rejected, on the side of meaning. According to a bilateral, two-sorted normative pragmatic theory, for the premises Γ to *imply* the conclusion A is for anyone who *accepts* all of Γ to be precluded thereby from entitlement to *deny* A—and in that sense, to be *implicitly* committed to *accept* A. According to a modal-mereological truth-maker semantic theory, for premises Γ to imply the conclusion A is for every fusion of truth-makers of all the propositions in Γ with falsity-makers of A to be *im*possible states.

On the basis of Hlobil's pragmatic-semantic isomorphism result, I invited us to think of these two very different kinds of discursive metavocabulary as offering a stereoscopic view of one single topic: reason relations. The *very same* relations of implication and incompatibility can be understood *both* in deontic normative terms of which constellations of doxastic commitments (acceptances and rejections) are *inappropriate*, in the sense that interlocutors cannot be jointly entitled to all of them, *and* in alethic modal terms of which fusions of sets of states (truth-makers and falsity-makers) are *impossible*. This is ***bimodal conceptual realism*** about reason relations.

This thesis, and the pragmatic/semantic isomorphism that justifies it, articulate an understanding of the bilateral distinction between practical doxastic attitudes of acceptance and rejection and the bivalent distinction between semantic values of truth and falsity as manifestations of a single *basic discursive bipolarity*. After all, doxastic *acceptance* of a claimable content, of the sort paradigmatically manifested in speech acts of *assertion*, is practically taking or treating that content as *true*. And doxastic *rejection* of a claimable content, of the sort paradigmatically manifested in speech acts of *denial*, is practically taking or treating that content as *false*. These relations at the level of sentences between pragmatic attitudes and semantic values are not controversial. What *is* controversial is whether either a pragmatics-first or a semantics-first order of explication is possible here: whether one or the other set of opposed metaconcepts can be made sense of sufficiently independently of the other to support an explication of true and false in terms of acceptance and rejection, or the other way around.

I want to start off today by suggesting that we will not adequately understand these two manifestations of this basic discursive bipolarity if we remain at the level of *sentences*. That same bipolarity necessarily shows up also at the level of reason relations among sentences. One form it takes there is the distinction between relations of *implication* and relations of *incompatibility*.[17] In the context of my larger argument, this claim should not be surprising. For the line of thought I have been developing locates the structure common to bi*lateral pragmatic* metavocabularies and bi*valent semantic* metavocabularies at this higher, suprasentential, reason-relational level.

The structural and functional connections between the sentential true/false and accept/reject dyads, on the one hand, and the reason-relational implication/incompatibility dyad, on the other, are subtle, and not obvious, however. We can begin to make those connections visible by thinking about the relations between classical bivalent truth values that make it possible for them to play their appointed role in semantic theory. First, there are two of them;

[17] In the third lecture, we will be much concerned with another version of the bipolarity that shows up within implication relations: the distinction between the role sentences play as *premises* of implications and the role they play as *conclusions* of implications. This version is already implicit in the bilateral pragmatic understanding of reason relations.

they are *different* truth values. But they are not *merely* different, they are *exclusively* different. *True* and *false* function like incompatible properties such as *square* and *triangular*, not like compatible, merely different ones such as *square* and *red*. Late in the game we might consider relaxing the prohibition on anything being both true and false—and in my third lecture I'll have something to say about logics that do that. But the original bivalent conception forbids overlap of extension between true and false propositions. Fine's sophisticated truthmaker semantics retains this basic structural feature, via his Exclusivity condition: any fusion of truth-makers of sentence with falsity-makers of that same sentence must be an impossible state. That expresses his modal reading of the incompatibility of truth and falsity. On the pragmatic side, there is a corresponding normative prohibition on accepting and rejecting the same claim: one can never be entitled to such commitments.[18] There is a deep connection between this shared exclusivity feature of truth values and doxastic attitudes, construed in the one case in alethic modal terms and in the other in deontic normative terms, on the one hand, and the reason relation of incompatibility of claimables, on the other hand.

The exclusiveness or incompatibility of both paired truth values and paired doxastic attitudes is a *symmetric* relation among them. It does not privilege truth over falsity or acceptance over rejection. But there is also an important *a*symmetry between them. We cannot systematically swap falsity for truth and truth for falsity while preserving the applicability of semantic theory. Truth and falsity are not just *different* and *incompatible* semantic properties, truth is in some sense the *good* one, the one we *want* or prefer. And something analogous is true on the side of pragmatics. Practical attitudes of acceptance and its manifestation in speech acts of assertion have a certain pragmatic priority over attitudes of rejection manifested in speech acts of denial. While it might be more difficult to articulate just how acceptance is pragmatically privileged relative to rejection, it is intelligible as a priority of *taking*-true over *taking*-false.

The key claim I want to make here is that the sense in which truth is *primus inter pares* relative to falsity in the classical bivalent understanding of truth values is that it is what is

[18] This corresponds to Reflexivity (RE), that each claimable implies itself. The expressive completeness result for the logic NMMS will depend on its generalization, Containment (CO).

preserved by good implications. The reason one cannot systematically swap falsity for truth while preserving a functioning semantic theory is that falsity-preservation does not yield a usable notion of implication.[19] This diagnosis of the nature of the asymmetry between truth and falsity as deriving from their role in defining the goodness of implications is supported by the division of labor in multivalued logics, which split the functions of bivalent truth values into two parts. The symmetric exclusiveness or incompatibility of truth values takes the form of a variety of *multivalues*—in three-valued logics, adding a third value to *true* and *false*. The distinction among multivalues between those that are preserved by good implications and those that do not play that role is then expressed by *designating* some of the multivalues, where to say that a multivalue (paradigmatically, *true*) is *designated* is to say that the good implications are just those that do *not* have premises all of whose multivalues are designated and a conclusion that is *not* designated.[20]

This way of looking at things explains how the distinction among reason relations between incompatibility and implication can be seen to be a species at a higher level of the same genus of basic discursive bipolarity as the distinctions at the level of sentences between truth and falsity, on the semantic side, and between acceptance and rejection, on the pragmatic side. Incompatibility relations capture the modally robust symmetric exclusion or repulsion aspect of the basic bipolarity, which is common to the semantic and pragmatic oppositions between true/false and accept/reject. In the classic setting (though not in the one we eventually recommend), implication relations explain the nonsymmetric privileging of some sentential truth values over others: truth (or the designatedness of multivalues) is the property of sentences preserved by good implications. Further, I think this line of thought gives us reason to think of the version of the dichotomy at the level of reason relations as in an important sense more fundamental than the versions that show up at the sentential level. For this account articulates the *meta*metaconceptual relational features of bivalent truth values that explain their metaconceptual capacity to express important properties of sentences. I originally introduced the concept of reason relations of implication and incompatibility as normatively governing which doxastic commitments count as a reasons *for* which others, and which count as reasons *against* which others. In the first

[19] We can formulate such a notion in terms of falsity, requiring that any implication that has a false conclusion must have at least one false premise, but that requires changing the classical model of good implication in terms of *preservation* of some semantic value. In fact, our substructural motives require us in any case to move beyond this model, since it builds in transitivity of consequence.

[20] Swapping designated for designated values never changes the goodness of an implication, and swapping two sentences with the same multivalue never changes the designatedness value of any compound in which they occur as components. Lecture III exploits the idea of substitutions that preserve the goodness of implications rather than truth, assimilating sentences insofar as they are intersubstitutable *salva consequentia* rather than *salva veritate*.

instance, these are reasons to *accept* and reasons to *reject*, forwarded in justificatory defenses of and critical challenges to doxastic commitments.

I have argued that we will not properly understand the semantic and pragmatic properties of sentences that exhibit the basic discursive bipolarity of true/false and acceptance/rejection unless the story includes reason relations of implication and incompatibility. And we saw last time that there is a robust bimodal isomorphism between pragmatics and semantics that also holds at the level of reason relations. Together, these considerations give us reason to look to reason relations in order to understand the truth-evaluable, acceptable/rejectable contents expressed by declarative sentences, which are what stand to one another in the relations of implication and incompatibility that are amphibious between what is specified in the bilateral pragmatic and truth-maker semantic metavocabularies, each of which exhibits its own version of the basic discursive bipolarity. On the pragmatic side, it is constellations of doxastic commitments to accept/reject that stand in relations of implication and incompatibility, and on the semantic side, it is Finean worldly propositions: pairs of sets of truth-making and falsity-making states satisfying Exclusivity.

Because the very same reason relations can hold in both the pragmatic and semantic settings, we can consider what kind of thing stands in those relations, just insofar as those relations are shared. This will be the rational or conceptual aspect of the propositions that can be specified in the two quite different ways. Because they stand in reason relations, I will call the relata '*rational* propositions' when they are specified in terms of their role in reason relations alone. They are functional roles those very different kinds of items can play with respect to the (potentially) shared implications and incompatibilities. We need not go so far as to identify truth-evaluable sayables and claimables with those roles in reason relations in order to investigate the rational dimension of such propositional contents, in the sense of the aspect or dimension of content that consists in playing that role in reason relations. We can call that specifically 'conceptual' content.

In my third lecture, I introduce a pure formal model-theoretic semantics of propositional conceptual roles in this sense: implication-space semantics. It is 'pure' in that the *only* resources it draws on are the reason relations items in a lexicon of sentences stand in to one another,

according to what in the first lecture I called a 'vocabulary.' A vocabulary, in that sense, is a relational structure that consists of a domain and a set of relations on that domain. In the basic case the domain is a *lexicon* of sentences. The relations are reason relations among the sentences, in the form of a set of pairs of sets of sentences. If the lexicon is *L* and the pair of two subsets of *L*, *S* and *S'* is included in the reason relations of the vocabulary, that means that the implication whose premises are the sentences of *S* and whose conclusions are the sentences of *S'* is a good implication, according to that vocabulary. (Incompatibility relations are encoded by marking the incoherence of sets of sentences by pairing them with the empty set.) So construed, vocabularies abstract away from the bilateral pragmatic and truth-maker semantic accounts of what *makes* implications good: the preclusion of joint entitlement to a set of doxastic attitudes and the impossibility of the state resulting from fusing truth-making and falsity-making states, respectively. The vocabulary just specifies which consequences among sentences hold, ignoring *why* or in what sense they do. Remarkably, these spare raw materials suffice for an expressively powerful and flexible formal model-theoretic representation of the conceptual roles played by declarative sentences.

II. The Structure of Reason Relations

The top-down order of explication I am pursuing, which appeals to relations of implication and incompatibility in order to understand the propositional contents expressed by declarative sentences, faces clarificatory demands concerning the structure of reason relations that traditional bottom-up orders of explication do not. In Fine's truth-maker semantics, for instance, the structure of consequence relations is determined by and inherited from the modal and mereological structure of propositions, just as in earlier views consequence relations were read off from the truth conditions of the sentences that show up as premises and conclusions. What considerations specifiable at the level of reason relations put structural constraints on implications and incompatibilities? We have so far seen two. To begin with, reason relations come in two flavors, corresponding pragmatically to reasons for and reasons against—reasons entitling interlocutors to accept and reasons entitling interlocutors to reject—doxastic commitments. I have argued that we should think of this bit of structure as the manifestation at the level of reason relations of the same basic discursive bipolarity that shows up in traditional semantics as the distinction between truth and falsity, and in bilateral pragmatics as the distinction between acceptance and rejection. I claimed, though I did not argue, that there is a further structural element to the distinction between the two flavors of reason relation: incompatibility is symmetric, while implication is nonsymmetric.[21]

What further structural restrictions might the metatheoretic role envisaged for reason relations impose on the consequence relation? We want to assume nothing about the structure of the bearers, sets of which stand in reason relations to each other. On our spare conception of a vocabulary, implication relations are just sets of pairs of sets of sentences, each such pair being thought of as pairing a set of premises and a set of conclusions that follow from those premises in the sense of 'follows' being represented. What considerations could be appealed to in imputing a structure to consequence relations as such?

[21] ROLE-contributor Ryan Simonelli makes an ingenious and compelling pragmatic social Dutch Book argument for the necessary symmetry of incompatibility in "Why Must Incompatibility Be Symmetric?" *The Philosophical Quarterly*, Volume 74, Issue 2, April 2024, Pages 658–682, https://doi.org/10.1093/pq/pqad078.

From the point of view of a top-down, implication-relations-to-sentential-contents order of explication, it is perhaps surprising that the logistical tradition has a well-defined, widely agreed-upon answer to this question, even though that tradition is not at all pursuing the project that motivates me to ask it here. A century or so ago, Tarski and Gentzen, the founders respectively of the model-theoretic and proof-theoretic traditions in logic, put forward basically the same proposal for necessary and sufficient structural conditions for relations among sets of sentences to qualify as genuine *consequence* relations.[22] Their official topic was specifically *logical* consequence relations. But unlike his, their characterization of the required structure did not appeal to the internal logical structure of the sentences that stand in those consequence relations. Indeed, Tarski sometimes omits the qualification 'logical' and refers to his topic just as 'consequence relations.'

Tarski's view is that what is essential about consequence relations is that they correspond to topological closure operators. In our terms, he takes a vocabulary to be a lexicon L of sentences, and a two-place relation between sets of sentences pairing each subset $X \subseteq L$ with its consequence-set Cn(X), satisfying these conditions:

Containment (CO): $X \subseteq Cn(X)$.

Monotonicity (MO): $X \subseteq Y \Rightarrow Cn(X) \subseteq Cn(Y)$.

Idempotence (CT): $Cn(Cn(X)) = Cn(X)$.

(These are variants of the Kuratowski axioms for topological closure operators.)

It will be helpful to think of these principles in other notations. We can write '$\Gamma|\sim A$' to say that the premise-set Γ (a subset of the lexicon L) *implies* the sentence A. The *explicit* content of the premise-set Γ consists of the sentences that are elements of that premise set. The *implicit* content of the premise-set Γ is, in a very literal sense, whatever it implies. Put in these terms, the structural principle of Containment says that all of the *explicit* content of every premise-set is also part of its *implicit* content. Monotonicity says that *adding* to the explicit content of a premise-set never *subtracts* anything from its implicit content. Idempotence says that making

[22] Alfred Tarski "On Some Fundamental Concepts of Metamathematics" (1928), in J.H. Woodger (trans.) *Logic, Semantics, and Metamathematics: Papers from 1923 to 1938 by Alfred Tarski* [Oxford University Press 1956] Ch. III, pp. 30-38. Gerhard Gentzen, "Investigations into Logical Deduction" *American Philosophical Quarterly* Volume 1, Number 4, October 1964, pp. 288-306.

implicit content explicit never *adds* any implicit content. Together they guarantee that drawing consequences from a premise-set is a cumulative enterprise that leads path-independently to a single, stable conclusion set: the rational closure of the original premise-set, the set of all its consequences.

These are all principles concerning the structure of implication relations. They address only the role sentences play as premises and conclusions of implications, regardless of what internal structure those sentences might or might not have. The closure principles articulate a structure common to the consequence relations of traditional logics—paradigmatically, classical, modal, and intuitionist logics—and they hold of the reasoning in mathematical proofs. Those virtues are sufficient to confer some plausibility on the claim that the topological closure conditions specify necessary and sufficient conditions for a binary relation between sets of sentences to qualify as a genuine *consequence* relation.

There are at least two principled reasons one might accept such a definition. First, one might take it that 'consequence' just *means* specifically *logical* consequence. This is the claim that genuine reason relations are always, at base, *logical* reason relations: implication is logical deducibility and incompatibility is inconsistency. Behind such a definition is ***logicism* about reasons**: the view that in the end, all *good* reasons (whether for or against) must always be *logically* good reasons. The *Tractatus* is the purest example of such an account. I think not many contemporary philosophers would defend this sort of universal logicism about reasons (even as properly restricted to doxastic rather than practical reasons). A weaker, fallback position is what could be called **'*structural*' logicism about reasons**. It acknowledges that some good implications might not be logically good, but insists that even *material*, *non*logical consequence relations must share the topological closure *structure* of logical consequence relations, on pain of not qualifying as *consequence* relations, in the sense that has normative significance for reasoning practices.

In spite of the considerable weight of authority and tradition behind treating topological closure structure as a necessary condition for genuine consequence relations, the stubborn fact is that outside of the artificial formal languages of logic and mathematics, a lot of actual reasoning

conducted in natural languages is *defeasible*, rather than *monotonic*. Premise-set Γ can imply A, even though if we add further sentences to Γ, the result no longer implies A. This sort of defeasibility of evidence, its status as merely *probative* rather than *dispositive* in support of a conclusion, is ubiquitous in practice—even in the most institutionalized contexts of reasoning, such as law and medicine. Nonmonotonicity is also a familiar feature of probabilistic reasoning, where new information can change the relevant reference-class with respect to which frequencies are assessed. Nor do speakers cite defeasible reasons just for convenience, because it would be too tedious to include all the provisos and conditions needed to render the implication indefeasible. There might be no systematic characterization of all the necessary qualifications. *Ceteris paribus* clauses should be understood as explicitly acknowledging the existence of unspecified defeasors, rather than as somehow turning a defeasible reason into an indefeasible one. (A Latin phrase whose utterance could do that is called a 'magic spell.') And although it is less commonly remarked, incompatibility relations are no less defeasible in general than implications.[23] The view that defeasible reasons, for or against, must be elliptical for 'full' reasons that are not defeasible is a consequence of commitment to a structural logicism that is controvened by actual reasoning practices.

The response of some philosophers to the empirical prevalence of nonmonotonic nonlogical implications has been to develop nonmonotonic logics. I think that is a mistake—and not just when it is justified by commitment to an objectionable logicism or structural logicism about reasons in general. The proper task should be understood rather to be to construct a logic adequate to *express* nonmonotonic implication (and incompatibility) relations. The distinction between the two enterprises is subtle, but important. In order to make it properly visible, I will need to motivate ***logical expressivism***. That is the claim that:

> The expressive task distinctive of *logical vocabulary* as such is to *make reason relations* of implication and incompatibility *explicit* in the form of claimable propositional contents of declarative sentences.

[23] This fact severely constrains the range of applicability of approaches to nonmonotonic reasoning that essentially depend on classical notions of inconsistency, such as default logics, e.g. Reiter, Raymond, 1980. A Logic for Default Reasoning. *Artifical Intelligence*, 13: 81–132.

That is where I am heading. First I want to argue that focusing on monotonicity is already making a kind of mistake, by overlooking other structural conditions on reason relations that are more interesting and important.

Monotonicity of implication is a very strong, doubly quantified *weakening* principle. (Adding premises to an implication is weakening it: it is a stronger claim that the conclusion follows from a proper subset of the premises.) MO says that for *any* arbitrary good *implication* Γ|~A and *any* arbitrary further *sentence* B, Γ,B|~A is guaranteed also to be a good implication if Γ|~A is. If we take seriously the idea that MO is *too* strong a constraint on rational consequence relations generally, there are accordingly two universal quantifiers that might be restricted. We can consider principles that only allow weakening *of* certain kinds of *implications*, but not all, and we can consider principles that allow weakening only *with* certain kinds of *sentences*, but not all. Containment, CO, is a restricted monotonicity principle of the first kind. It focuses on just one class of implications, those licensed by Reflexivity (RE), which says that all implications of the form A|~A are good. CO says that all implications of that form can be arbitrarily weakened by any sets of additional premises whatsoever. In the idiom I have suggested, CO says that any implication is good whose conclusion, what the implication certifies as part of the *implicit* content of the premise-set, is already part of the *explicit* content of the premise-set, that is, is one of the premises. This is an antecedently plausible constraint to put on a conception of "following from." And, as we shall see, it turns out to cut at important joints.[24]

The principle called *Cautious Monotonicity* (CM), by contrast, quantifies over all good implications, while restricting what can be added to the premise-set, by licensing only weakenings that add sentences meeting a special condition. That condition is defined relative to the original premise-set. The plausible idea is that while there might be *some* sentences whose addition to a premise-set infirms the implication of some of its consequences, it is safe to add to the premise-set sentences that are already implied by it. Here is a Gentzen-style sequent calculus formulation of this rule, which should be read as saying that if all the implications above the horizontal line are good, then so is the implication below the line:

[24] Classical logic codifies just the sequents that follow from all CO-instances. We'll introduce a modal operator that marks off local regions of classicality in this sense.

Cautious Monotonicity (CM): $\frac{\Gamma|\sim A \quad \Gamma|\sim B}{\Gamma,B\,|\sim A.}$

If some sentence B is already part of the *implicit* content of premise-set Γ, then adding it to Γ as an *explicit* part of the premise-set does not subtract any implicit content. Anything that followed from Γ also follows from Γ together with any of its other consequences, since they are already implicitly contained in it.[25]

Cautious Monotonicity brings into view an important operation on implications. For it involves comparing the consequences of two premise-sets, Γ and Γ together with another sentence that is part of its implicit content. In effect, we are looking at the effect of *explicitation*, in the sense of making some of the *implicit* content of a premise-set *explicit*, by adding that consequence to the premise-set. This is moving a sentence from the right-hand, conclusion side of the turnstile, to the left-hand, premise side. Explicitation in this clear structural sense is a relation between implications (that is, a relation between reason relations). It is important because on the pragmatic side it normatively governs the process of *inferring*, understood as *acknowledging*, as an *explicit* (avowed) commitment, something that one was only *implicitly* committed to, in the sense that it was implied by one's other commitments.[26] Doing this is a central, crucial form of rational inferential activity: extracting the rational consequences of one's beliefs. So it is worth thinking a bit about the relations between explicitation and structural restrictions on reason relations.

In that connection, CM says that explicitation never *loses* implicit content: anything implied by a premise-set is also implied by that premise-set together with any of its other consequences. It has a dual, which says that explicitation never *adds* implicit content: anything implied by a premise-set together with some of its consequences is already implied by the premise-set alone. This structural principle is

---

[25] Just as there is an analogue of MO for incompatibility, which says that if premise-set Γ is incoherent (so any premise in Γ is incompatible with the remainder of Γ), then so are all its supersets, there is also an analogue of CM. It says that if Γ is incoherent, so is any superset of it that results from adding only consequences of Γ.

[26] Though related, this sense of 'implicit', and of inference as moving from the implicit to the explicit, is different from the one mentioned last time, where preclusion from entitlement to accept a claim implicitly commits one to rejecting it, and *vice versa*. For explicitation in the sense discussed here is not defined in terms of the basic discursive bipolarity of accepting/rejecting—even if we ultimately understand reason relations in terms of that bipolarity.

Cumulative Transitivity (CT): $\frac{\Gamma, B \mid\sim A \quad \Gamma \mid\sim B}{\Gamma \mid\sim A.}$

This is just a sequent-calculus version of the Tarskian transitivity-as-idempotence closure principle, since it says that consequences of consequences are already themselves consequences of the original premise-set. It, too, governs the paradigmatically rational pragmatic process or practice of inferring, in the sense of acknowledging consequences of one's commitments.

The connection to rational explicitation and the symmetric duality of CM and CT that consists in their just being re-arrangements of the same three sequents are what I meant by suggesting that the pairing of MO with CT is both less important and less natural than appealing to the dual explicitation principles, quite apart from the empirical observation that nonlogical implications are not always monotonic. Although CM and CT do not require implication to be a topological closure relation, they do define a weaker, but still significant kind of rational equilibrium. Since CM says that explicitation never *subtracts* consequences from a premise-set, and CT says that explicitation never *adds* consequences to it, together they entail that ***explicitation* is *inconsequential***. Making the implicit, consequential content of a set of premises explicit as further premises never changes the implicit content, what follows from those premises. All the premise-sets that result from any given one by adding some of the sentences it implies have exactly the same consequence sets. In this sense, Cautious Monotonicity and Cumulative Transitivity define a structural condition on reason relations that we can call *explicitation closure*. Consequence relations that are closed under explicitation form a natural kind. It includes the fully monotonic, topologically closed implication relations, but also many *non*monotonic ones. Explicitation-closed implication relations play a special role in the paradigmatically rational inferential process of discovering and acknowledging explicitly the implicit consequences of one's commitments.

Are there any relations that deserve to be thought of as *consequence* relations that are *not* closed under explicitation? Yes. Explicitation is *not* always inconsequential. Here is one very general, practical process governed by an implication relation where changing the status of a claimable from conclusion to premise is of considerable significance. Consider a database at an experimental physics installation such as a superconducting supercollider. All the observational

data from the supercollider is put into a database. An inference engine in the form of a scientific theory is then clamped onto the database, to extract consequences of those observations according to the theory. Many of those consequences will themselves be sentences expressing observables. If observation qualifies one of those theoretically predicted consequences to be placed in the database of observations, that experimental confirmation of the theory can be an event of great import. Confirming one prediction might well offer reasons for making further predictions and changing confidence in or even endorsement of others, for the predictions need not in general all be compatible with one another. In this context, explicitation, as changing the status of a sentence from expressing a conclusion to expressing a premise, has the significance of empirical confirmation of the theory used to make the prediction. Far from being inconsequential, this sort of explicitation is at the core of scientific practice and its constraint by empirical observation. Of course, not all consequence relations attach this sort of additional significance to the distinction between premises and conclusions. But the point is that they *can*, and when they *do*, CM and CT need not hold. At the least, we have good reason to want to make sense of consequence relations that are *hyper*nonmonotonic, in not even being cautiously monotonic or cumulatively transitive.

Looking at the structure of implication relations from the vantage point of explicitation is considering the constraints that stem from relations among implications, for that is what explication is. We care about that explicitation relation because it normatively governs attempts to extract consequences from a set of commitments. An *explicitation path* from a premise-set Γ is a sequence of supersets of Γ, each one resulting from the previous one by its adding as a premise some consequences of the previous premise-set. We are used to thinking of explicitation as inevitably leading to a foregone conclusion: the same one no matter what consequences we acknowledge first, and which later. The rational closure of a set of premises Γ comprises *all* of the consequences of Γ. But there is such a thing as the rational closure of a set of commitments only in the special case where CM and CT hold, so that consequence is explicitation closed. In hypernonmonotonic cases, where CM and CT are not guaranteed to hold, explicitation paths can diverge. Which conclusions one can reach from the same base Γ depends on the order in which one extracts consequences from it. Some that are passed over early can become inaccessible. This is rational *hysteresis*, or path-dependence. The process of inferring, in the sense of

following out an explicitation path, is, in an explicitation-open setting, an essentially *historical* one that can lead far afield from its starting point.

People disagree about whether believing all the consequences of one's beliefs is an epistemic ideal. The usual objection is that it is impossible for us poor finite, forked creatures to do that (the consequence-set is too large, some of the conclusions too distant), and ought implies can. But both parties to that dispute think there is something definite to *mean* by "all the consequences of one's beliefs." That is what they argue can or cannot and ought or ought not to be believed. My point is that whether it does make sense depends on the structure of the reason relations involved. There is a specifiable boundary between consequence relations for which it does, and consequence relations for which it does not, so much as make sense. That boundary is explicitation closure.

There are structural relations between reason relations other than those on the spectrum I have been describing, which I have not discussed here, such as the principle of *explosion*. It connects incompatibility and implication relations by dictating that incompatible premise-sets have the whole lexicon as their consequence set. These are all *rational* structures, rather than *logical* or *semantic* ones. As here described, none of these structural relations among implications and incompatibilities should be thought of as articulating *logical* structure. For none of their specifications appealed to the appearance of any logical vocabulary in the sentences that stand in the reason relations that are structurally related to others. Indeed, no appeal was made to any *semantic* properties of those sentences either. The top-down order of explication being pursued here would have us understand structures of reason relations as *affecting* rather than *reflecting* the propositional contents expressed by the sentences that play roles as premises and conclusions. This underlying *rational* level, the level of reason relations, is prior both to logic and to a semantics of sentential propositional conceptual contents according to the order of explication I am pursuing.

I want to claim that even the most structurally relaxed of these kinds can serve the basic pragmatic function of determining reasons for and against to serve in defenses and challenges of claimings, where a principal criterion of adequacy of doing that is underwriting the isomorphism with a truth-maker semantics, at the level of what therefore display at least those credentials for being called *reason* relations. But I want to look more closely at two further metatheoretic offices that the concept of reason relations is called upon to carry out. For the pragmatic and semantic formal metavocabularies for talking about reason relations do set substantial,

reasonably definite criteria of adequacy on conceptions of the structure that characterizes *reason* relations as such. Can we do logic with radically open-structured relations of implication and incompatibility? Do hypernonmonotonic reason relations allow the definition of a tractable notion of the propositional conceptual contents expressed by declarative sentences in virtue of the role those sentences play as premises and conclusions of implications and as standing in incompatibility relations? I hope to show that both these roles can be played by structurally open reason relations as well as structurally closed ones. Reason relations of any structure that can clear those substantial hurdles in logic and semantics will be able to play their role in pragmatics, as normatively governing a minimal discursive practice of making, challenging, and defending claims. And those open (both topologically and explicationally) reason relations will suffice to correlate, up to isomorphism, the use of sentences described in such a pragmatics with a truth-maker modal-mereological account of their meanings. In my third lecture, I will introduce a substructure-tolerant implication-space semantics that defines and manipulates pure propositional conceptual roles defined solely in terms of reason relations. I'll turn now to the issue of logic, and introduce implication-space semantics as a pure theory of propositional conceptual roles next time.

III. Logical Expressivism

What might be called the "reasons question" in the philosophy of logic is how logic is related to reasons and reasoning. My first claim is that the beginning of wisdom in addressing that question is to learn from Harman's point that the relation of logic to reasoning *practices* is mediated by reason *relations* of implication and incompatibility. Earlier, I mentioned the *logicist* answer to that reasons question: all good reasons are at base *logically* good reasons. I have just been arguing against the weaker thesis of *structural* logicism, which maintains that all nonlogical consequence relations must share the closure structure of *logical* consequence relations, on pain of not qualifying as relations of *rational* implication. Since logicism about reasons entails structural logicism, if there are structurally open reason relations, logicism cannot offer a general answer to the reasons question.

If logic does not determine in general what implications and incompatibilities hold among nonlogical sentences, how is it related to those sometimes open-structured reason relations? My second claim is that, properly understood, the task of logic is not to *determine* nonlogical reason relations, it is to *express* them. Put slightly more carefully, the expressive task that distinguishes logical vocabulary is making implication and incompatibility relations among sentences of nonlogical base vocabularies explicit in vocabularies that have been extended from those bases by the addition of that logical vocabulary. This general form of response to the reasons question is ***rational expressivism*** about logic. It is a kind of 'expressivism' because it understands the defining task of logic to be a matter of what it makes it possible for its users to *say*, rather than, for instance, anything about what it makes it possible for them to *prove*. It is a 'rational' expressivism because what is expressed is taken to be reason relations, rather than, for instance, some kind of attitude (as in traditional metaethical expressivism).[27]

[27] Which has been ingeniously revived and updated to apply to this sort of case by Luca Incurvati and Julian Schlöder in *Reasoning with Attitudes: Foundations and Applications of Inferential Expressivism* [Oxford University Press, 2023].

In my first lecture I sketched a pragmatic account of a minimal discursive practice in which participants use sentences to assert and deny claimables, which essentially involves rationally defending and challenging claimings with reasons for and against them, which are in turn determined by reason relations of implication and incompatibility among them. By doing what they do, such users of a base vocabulary accordingly practically acknowledge the reason relations that normatively govern their giving and asking for reasons. But they need not be able to *say* what sentences imply or are incompatible with each other. Logical vocabulary gives them the additional expressive power to do that. The idea is that the conditional "If A then B" says that A implies B. Adding negation then makes "If A then *not*-B" available to express the incompatibility of A and B. Logical vocabulary makes it possible to make reason relations explicit, in the sense of *sayable*, assertible and deniable, rationally challengeable and defensible declarative sentences. The defining expressive function of conditionals and negations is to codify reason relations in the form of rational propositions, that is, in the form of claimables that can themselves stand in reason relations of implication and incompatibility—both to sentences of the lexicon of the prelogical base vocabulary and to other logically complex sentences formed from them.

We can make this way of thinking about the expressive role of logic more precise by formulating it in terms of the simple relational structures that in my first lecture I called 'vocabularies.' As I am using the term, a vocabulary is an ordered pair of a *lexicon* and a set of *reason relations* defined on that lexicon. The lexicon is just a set of sentences (or other "bearers," such as Finean propositions). The reason relations can be thought of as the set of *good* implications. For technical reasons, we think of implications as pairs of a premise-set of sentences and a conclusion-set of sentences. (Incompatibilities are coded as implications with an empty conclusion-set.) On the expressivist account, a logic is a means to extend a base vocabulary to a supervocabulary of it, in the sense of a vocabulary whose lexicon is a superset of the base lexicon and whose reason relations contain those of the base vocabulary. To introduce logical locutions into a language, then, one must define a function that first expands the lexicon of any base vocabulary by adding new sentences that are logical compounds of old ones, and then computes the reason relations of the newly expanded lexicon from the reason relations that govern the use of the base lexicon.

The first, syntactic, part of this process is easy. If the base lexicon is a set of sentences L_{Base}, then the new, logically extended lexicon L is fully defined by:

$\mathbf{L_B \subseteq L}$ and
$\mathbf{A,B \in L} \Rightarrow$ $\mathbf{\neg A \in L}$ and
$\mathbf{A \rightarrow B \in L}$ and
$\mathbf{A \& B \in L}$ and
$\mathbf{A \vee B \in L}$.

The question is: How can we compute the implications and incompatibilities that govern this new, logically extended lexicon of logically complex sentences, entirely from the implications and incompatibilities that govern the old, base lexicon of logically atomic sentences? We can see how the logically extended lexicon can be elaborated or computed from the base lexicon. How are the new reason relations elaborated from or determined by the reason relations of the base vocabulary?

My third claim is that the ideal metavocabulary for specifying those relations is the sequent calculus that Gerhard Gentzen introduced in the founding document of proof theory. Its basic idea is to treat reason relations, specifically implications, as mathematical objects, called 'sequents.' It operates on those objects and permits the formulation of rules relating them. In short, the sequent calculus is an expressively powerful metavocabulary for specifying relations among reason relations.[28] Sequent rules always have the metainferential form: if *these* sequents (codifying implications) are good, then so are these others. The input sequents are written above a horizontal line called an 'inference line,' with the output sequents written below it. (Metainferences of this sort can be strung together, to yield derivations of some sequents from others.) Metainferential rules are of two kinds: structural and operational. The structural rules do not depend on anything about the lexical items involved in the sequents, except their identity or nonidentity to one another. We can formulate monotonicity this way as:

Monotonicity (MO):

$$\frac{\Gamma \mid\sim \Delta}{\Gamma, A \mid\sim \Delta.}$$

Meta-Inference Line

[28] It is a metavocabulary in that it is a metalanguage for discussing vocabularies. There is *also* a formulation of it *as* itself a vocabulary in my technical sense. But pursuing and justifying that idea is not part of the project of these lectures.

If the premise-set Γ implies the conclusion-set Δ, then so does Γ together with any arbitrary sentence A of the lexicon.

The operational rules include relations among implications that contribute to the meanings of logical locutions. One that is important to my story is the right rule for the conditional, which captures its expressive role as making implications explicit in the form of sentences of the logically extended object language:

Deduction-Detachment (DD): $\Gamma, A \mid\sim B, \Delta$ / $\Gamma \mid\sim A \rightarrow B, \Delta$. Bidirectional Inference Line

(Here the double horizontal line means that the metainference is being stipulated to hold in both directions.) Another operational metainferential rule captures the expressive role of negation, needed to express incompatibilities as logical inconsistencies:

Incoherence-Incompatibility (II): $\Gamma, A \mid\sim \Delta$ / $\Gamma \mid\sim \neg A, \Delta$. Bidirectional Inference Line

(This is the multisuccedent version of $\Gamma \mid\sim \neg A$ if and only if $\Gamma \# A$, i.e. $\Gamma, A \mid\sim$.)

The important point is that sequent rules are a special-purpose way of *constructing* or *computing* the reason relations of an extended, logical supervocabulary from the reason relations of a prior base vocabulary. They do that by codifying reason relations of (meta)implication that hold among the sequents that themselves express the reason relations of some object-language material subvocabulary.

Rational expressivism about logic—the view that what distinguishes and demarcates logical locutions as such is their expressive role in making reason relations of implication and incompatibility propositionally explicit—both puts constraints on the admissible ways of computing the reason relations of a logically extended vocabulary from the reason relations of a nonlogical base vocabulary and sets criteria (norms, *desiderata*) for assessment of the adequacy of particular sets of metalogical sequent rules. First, the logically extended vocabulary must be *elaborated from* the base vocabulary, in the sense that both the lexicon and the reason relations of the extended vocabulary must be computed from those of the base vocabulary. Second, the logically complex sentences of the extended vocabulary must *explicate* the reason relations of the base vocabulary, as well as those of the logically extended vocabulary. A lot more will need

to be said to articulate and give a clear sense to this criterion of adequacy. But we already have the paradigm of the Deduction-Detachment rule that shows the sense in which the conditional makes it possible to form sentences that *say*, in the extended object language, *that* a particular implication holds. At this point I just want to observe that the explication condition entails a *conservativeness* requirement on the elaboration condition. Explicating reason relations, in the sense of making them explicit as the conceptual contents of sentences that themselves imply and are incompatible with others requires that doing so does not *change* what one is expressing. To achieve that, it is necessary and sufficient to require that in the logically extended vocabulary, all the implications and incompatibilities that involve only lexical items from the base vocabulary are just the same as they are in the base vocabulary from which the logical vocabulary is elaborated.[29]

I will say that a vocabulary that is conservatively ***el***aborated from and ***ex***plicative of some base vocabulary is 'LX' for that vocabulary. One measure of the rational expressive power of a logic, understood now as specified in a sequent-calculus metainferential language, is then determined by the variety of base vocabularies for which it *is* LX: from which the sequent rules conservatively elaborate it and whose reason relations it explicates (in a still-to-be-fully-explicated sense). The structural closure conditions on implication relations that I talked about earlier provide an appropriate scale on which to measure expressive power so understood. Some logics (as specified by metainferential sequent rules) fail to define conservative extensions of any nonmonotonic or nontransitive base consequence relations. That is true of Gentzen's canonical sequent specification of classical logic, his system LK. And even logics that could elaborate and explicate topologically open logics, which fail to satisfy the monotonicity principle MO, might well fail to be LX for logics that are not even explicitation-closed. The expressive ideal along this dimension is a logic that is *universally* LX: LX for *any and all* base vocabularies, regardless of the structure of their reason relations.

[29] This is a rationale, deriving from rational expressivism about logic, for the conservativeness requirement that Nuel Belnap introduced as a technical device to rule out 'tonkish' logical connectives (in "Tonk, Plonk, and Plink", *Analysis* 22 (6):130-134 (1962).

IV. A Logic for Open Reason Relations

I am happy to be in a position to share with you the good news that there is in fact such a universally LX logic. We call it *N*on*M*onotonic,*M*ulti-*S*uccedent logic, or NMMS for short.[30] [**The sequent-calculus connective definitions for this logic are in the Appendix.]** It has three remarkable properties. The first is that it is *expressively complete* in an unprecedentedly strong sense. Dan Kaplan showed how to associate each sequent or set of sequents whose premise-set and conclusion-set consist of logically complex sentences with a set of sequents in the base vocabulary, which relate only logically atomic sentences that occur in them, such that those sequents from the logical supervocabulary hold in all and only the NMMS-elaborations of bases in which just those atomic sequents hold. In this clear sense, the logically complex sequents *say that* the corresponding logically atomic sequents hold. For they are derivable just in case those reason relations hold in the base vocabulary. Fixing the lexicon of the base vocabularies, we can compute for any set of atomic sequents which logically complex sequents say that just those sequents hold, and for any set of sequents relating logically complex sentences we can compute just which base sequents they say hold. This is the precise version of the 'X' dimension of LX-ness: for *any* relations of implication and incompatibility that atomic sentences can stand in, NMMS permits the formulation of a single sequent in the logically elaborated supervocabulary that expresses just those ground-level reason relations, in the sense that that sequent is derivable just in case those reason relations hold in the base.[31]

This expressive completeness result is a more powerful version of the way that in the two-valued semantics for classical logic truth-tables let us go back and forth between the values of logically atomic sentences and those of logically complex ones. It applies to assessments of

[30] See the Appendix to this lecture for the connective definitions of NMMS. This logic was originally developed, and its expressive completeness proven, by Daniel Scott Kaplan, a member of our ROLE logic group, based on a single-succedent predecessor developed by Ulf Hlobil. It is discussed, and the results retailed here are proven in Chapter Three of *Reasons for Logic, Logic for Reasons*.

[31] There are some minimal conditions on the result, and niceties to be observed relating to Contraction. They are detailed in Chapter Three of *RLLR*. The fact that these expressive relations can be exploited in both directions depends on NMMS using only *reversible*—double inference line—sequent definitions of logical connectives. Gentzen's sort-of student Oiva Ketonen produced the first set of connective definitions in the sequent calculus that had this property, and so collapsed the distinction between derivability and admissibility.

candidate implications as good or bad, rather than assessments of sentences as true or false. (Next time we will look at the step from these extensional semantic associations to intensional ones, with the move from truth *values* to truth *conditions* paralleled on the inferential side by the move from the goodness of implications to their ranges of subjunctive robustness.) It is more powerful in that every finite set of atomic implications in the base vocabulary is represented by a *single* sequent in the logically extended consequence relation, and every implication in the logically extended consequence relation represents a *unique* set of atomic implications.

The second remarkable property of NMMS is that it is fully tolerant of open-structured or radically substructural base vocabularies. This feature has to do with both the *elaboration* dimension of LX-ness and the *explication* dimension. NMMS conservatively extends logically atomic base vocabularies that are nonmonotonic and nontransitive, those in which Cautious Monotonicity fails, and even those for which Containment fails, and its conditional and negated sentences codify the reason relations of such substructural reason relations. This is a substantial achievement, because it is easy for metainferential rules for connectives to enforce global structure, for instance monotonicity. An obvious example is a left rule for conjunction that says that if (in some context) a premise A implies something, then that same conclusion follows from the conjunction A&B, which rules out the addition of further premises infirming or defeating an implication.[32] Within very wide limits, the full expressive completeness result still holds for the elaboration of substructural base vocabularies by NMMS. So NMMS is provably *universally* LX: it elaborates any base vocabulary, no matter its structure, and does so in a way that is expressively complete, providing logical codifications of arbitrary collections of atomic reason relations.

The third remarkable property of NMMS is that it is essentially just classical logic. Understanding the sense in which NMMS *is* classical logic, and also what nevertheless distinguishes them, sheds significant light on the interrelations between the topologically and

[32] (The example depends on the assumption that the expressive function characteristic of conjunction is to make explicit the comma on the left of the turnstile. Relevance logic rejects this assumption.) Securing the desired tolerance of open-structured base vocabularies requires departing from usual sequent-calculus practice and mixing additive with multiplicative clauses in the rules for a single connective. That is usually avoided because it causes difficulties for the proof of Gentzen's Cut-elimination Hauptsatz. This is not a problem in our setting, since we do not want the global admissibility of Cut, which would be nonconservative over nontransitive base vocabularies.

explicitationally open structure of material, nonlogical reason relations and the expressive task distinctive of logical locutions. The first data point in understanding the intimate relations between classical logic and NMMS is this. In the fully topologically closed setting defined by Gentzen's full set of structural rules, NMMS yields exactly the same logically valid sequents as Gentzen's sequent-calculus version of classical logic, LK. In this sense, there are *many* equivalent sequent-calculus formulations of classical logic: different ways of defining the logical connectives that yield the same logically good sequents. As soon as we relax the structural requirement of topological closure, however, those hitherto equivalent formulations come apart. Differences among them that don't matter in closed structures turn out to make a difference in open-structured settings.

The second data point in understanding the intimate relations between classical logic and NMMS is that NMMS is *supraclassical* when applied to any base vocabularies that include all instances Containment, regardless of what other structural principles do or do not hold. That is, all classically valid implications and incompatibilities are included in the NMMS elaboration of every base vocabulary that includes as good all sequents where some premise is included in the conclusion-set. This takes us a step beyond the first point, that if these CO instances are the *only* good implications in the base vocabulary, then NMMS validates all *and only* classically valid reason relations.

So, out of all the variant ways of introducing the connectives of classical logic in the sequent calculus, NMMS is one that not only degrades gracefully, but continues to work fully, when we move from structurally closed to topologically and even explicitationally open settings. In those settings, NMMS still elaborates base vocabularies conservatively, it is supraclassical if those base vocabularies satisfy Containment, and yields exactly the classically valid implications and incompatibilities if it is applied to base vocabularies *all* of whose implications are instances of Containment. In all those open-structured and closed settings it remains *expressively complete* in the Kaplan sense. For every set of atomic base sequents, there is a unique logically complex sequent that is derivable just in case the base vocabulary to which it is applied contains those atomic sequents, and *vice versa*.

It is also possible to capture structural features of reason relations by moving beyond the propositional connectives of classical logic. By adding the right kind of modal operators to NMMS, we can define a logic with even more expressive power in substructural settings. Even in vocabularies where Monotonicity fails to hold as a global structural principle, because some implications or incompatibilities are defeasible by adding further premises, there can still be *some* indefeasible implications or incompatibilities. The otherwise serviceable implication from 'Tweety is a bird,' to 'Tweety can fly,' ceases to be good when the additional premise 'Tweety is a penguin,' is added to the premises. But 'Pedro is a donkey,' so 'Pedro is a mammal,' is not, or need not be, defeasible. Logical expressivism counsels that instead of imposing Draconian structural requirements such as monotonicity *globally*, that is, to *all* reason relations, we introduce logical locutions to mark explicitly *local* regions where some structural principles *do* hold, even though they are *not* assumed to hold *everywhere*. In our book, we show how to add this expressive power to NMMS. The idea is that one can define a monotonicity box to mark the *persistence* of the goodness of an implication under arbitrary addition of further premises, by a version of the principle that if not only does the premise-set Γ imply the conclusion-sentence A, but so do all the supersets of that premise-set, then Γ not only implies A, but also □A. Then we can use sentences marked with boxes to keep track of which implications hold persistently. In the same way, it turns out that we can explicitly express which implications are *classically closed*, in being not only monotonic but transitive.

V. Consequences and Conclusion

My aim here is to demonstrate the benefits from a change of perspective in thinking about logic. What matters most about a logic is not its theorems, nor its consequence or inconsistency relation: that is, the logical reason relations the logic gives rise to and enforces. What matters most is the *expressive power* it affords to make explicit a variety of material reason relations of *non*logical vocabularies. In our regimentation, particular implications and incompatibilities defined on nonlogical lexicons are codified as logically complex sentences in extensions of those vocabularies whose reason relations are defined by metainferential rules in a sequent-calculus metavocabulary. The right logic can make explicit the reason relations of *any and every* base vocabulary, even those with the most minimal structure. And the practical expressive capacities afforded by logical connectives even of less expressively powerful logics are transformative for language users.

We can compare two linguistic communities, one of which asserts and denies, and challenges and defends the resulting commitments according to the reason relations of a nonlogical base vocabulary with one that does the same with the lexicon and reason relations of the logically extended vocabulary computed from that base. The first group of speakers is *rational*. They can not just respond differentially to things, but can respond by making claims and judgments, and they can give and assess reasons for and against their claims and objections. But they can only disagree about, and critically assess the credentials of, doxastic commitments manifestable as assertions or denials. They cannot make claims about, or assess the credentials of the implications and incompatibilities they implicitly appeal to in their reasoning. They are *rational*, but not *self-consciously* rational. **Logic is at base an organ of rational self-consciousness**.

What confers that power is the full logically extended consequence and incompatibility relations, precisely because they relate logically complex sentences that do *not* codify reason relations that hold in virtue of logic alone: conditionals such as 'If it is raining, then the streets will be wet.' The purely logical reason relations articulate the contents of the logical locutions. And that is important. But the logically extended vocabulary articulates the contents of the nonlogical locutions of the base vocabulary. And that is *more* important.

We can define what it is for some implications and incompatibilities in the logically extended supervocabulary of a base vocabulary to hold 'in virtue of logic alone' as the ones that hold in the logical extension of every base vocabulary that meets some minimal structural condition. So, all the sequents of classical logic hold in every NMMS extension of bases in which all instances of CO (Containment) hold. Also, the implications and incompatibilities that hold in the NMMS extension of a base vocabulary whose reason relations consist *just* of the instances of CO (for its lexicon) are exactly those of classical logic. NMMS has been carefully sculpted to extend open-structured material consequence relations *conservatively*, not imposing monotonicity or transitivity on material bases that are not monotonic or transitive—indeed, even *hyper*nonmonotonic explicitationally open bases where Cautious Monotonicity fails. But the astonishing fact is that the purely *logical* part of any NMMS-extended base vocabulary is guaranteed to be fully structural and topologically closed: monotonic and transitive, so idempotent. In this strict and literal sense, NMMS is *not* a nonmonotonic logic. It is a monotonic, structurally classical logic for codifying nonmonotonic consequence relations—and nonmonotonic incompatibility relations, and nontransitive consequence relations, and ones for which even CM fails, and so on.

Looking back from the point of view of rational expressivism about logic—the thesis that the expressive role that distinguishes specifically *logical* vocabulary is making reason relations explicit—the logical tradition of the past hundred years or so can be seen to have made two large mistakes. The first is a version of the one made by the drunk who looks for his keys under the streetlamp rather than across the road where he dropped them, because the light is better there. In this application of the metaphor, the bright light is the clarity, perspicuity, and tractability of classical logic under the traditional bivalent semantic interpretation. The role of the keys is played by the codification of reason relations. Logicism about reasons is the view that reason relations *must* be codifiable by that spectacularly well-behaved logic. The second mistake is subtler. When logicians did take seriously the existence of substructural or open-structured reason relations, the form their efforts to codify them took was logics whose own reason relations were substructural. For if reason relations are at base logical reason relations, as logicism claims, then there must be open-structured *logical* reason relations behind those open-

structured material reason relations. That is, the motivation for nonmonotonic logics involves keeping logicism and giving up structural logicism. Expressivists do it the other way around. For we have discovered that we can split the difference by codifying radically open-structured reason relations with a logic, NMMS, whose own logical implication and incompatibility relations are topologically closed, just like those of classical logic. Further, that universally LX logic is supraclassical, and collapses to classical logic when applied to a flat prior, in the sense of a base vocabulary that consists only of CO instances.

My first lecture introduced the idea of reason relations of implication and incompatibility, and offered two sorts of explanations of them, one in bilateral pragmatic deontic normative terms and the other in truthmaker semantic alethic modal terms. One of my basic claims is that the minimal structural conditions on *rational* relations of consequence and incompatibility are *much* weaker than has traditionally been supposed. This time I addressed the question of whether, to what extent, and in what sense radically substructural 'reason relations'—whether merely topologically open or also explicitationally open—deserve to be thought of as genuine *reason* relations of implication and incompatibility. I have shown that they pass one test that is crucial for logical expressivists: they can be codified completely by well-behaved logical vocabulary—indeed, by what is in several real and important senses just *classical* logical vocabulary. The next task is to investigate what sort of conception of rational propositional content results from considering the role declarative sentences play in radically open-structured reason relations. Next time I will show how a full-blown implication-space formal semantics incorporating a tractable concept of proposition can be elaborated just from vocabularies, in the spare technical sense in which I have been using the term: a lexicon together with a set of pairs of sets of lexical items meeting the most minimal of conditions. The first criterion of adequacy that semantics satisfies is that it provides a sound and complete semantics for the universally LX logic NMMS. It turns out that that result can be generalized to many other logics, as well. Our ultimate interest, though, is in what implication-space semantics can teach us about the relations between meaning or propositional content and reason relations. That is my topic for next time.

End of Lecture II

Appendix to Lecture II:

Connective Rules of NMMS:

$L\neg$: $\dfrac{\Gamma|\sim\Delta,A}{\Gamma,\neg A|\sim\Delta}$ $R\neg$: $\dfrac{\Gamma,A|\sim\Delta}{\Gamma|\sim\Delta,\neg A}$

$L\rightarrow$: $\dfrac{\Gamma|\sim\Delta,A \quad B,\Gamma|\Delta \quad B,\Gamma|\sim\Delta,A}{\Gamma,A\rightarrow B|\sim\Delta}$ $R\rightarrow$: $\dfrac{\Gamma,A|\sim B,\Delta}{\Gamma|\sim A\rightarrow B,\Delta}$

$L\&$: $\dfrac{\Gamma,A,B|\sim\Delta}{\Gamma,A\&B|\sim\Delta}$ $R\&$: $\dfrac{\Gamma|\sim\Delta,A \quad \Gamma|\sim\Delta,B \quad \Gamma|\sim\Delta,A,B}{\Gamma|\sim\Delta,A\&B}$

$L\vee$: $\dfrac{\Gamma,A|\sim\Delta \quad \Gamma,B|\sim\Delta \quad \Gamma,A,B|\sim\Delta}{\Gamma,A\vee B|\sim\Delta}$ $R\vee$: $\dfrac{\Gamma|\sim\Delta,A,B}{\Gamma|\sim\Delta,A\vee B}$

Abstract of Lecture 3:

Roles and Reasons

For semantic inferentialists, the division of candidate implications into good and bad has a significance parallel to that of the traditional division of sentences into true and false. The move from truth values to truth conditions (extension to intension) is paralleled by the inferentialist's associating each implication with its *range of subjunctive robustness*: the premises and conclusions that can be added to it to make it good (if it is not) and to keep it good (if it is). Assimilating and ordering implications according to inclusions among their ranges of subjunctive robustness makes it possible to interpret sentences by pairs of the conceptual roles they play as premises and as conclusions in implications, both good and bad. Operations on such conceptual roles make it possible to formulate semantic connective definitions that are sound and complete for the maximally expressively powerful logic NMMS introduced in the previous lecture. That procedure turns out to generalize to all connective definitions formulated in a suitable sequent-calculus metavocabulary. Any constellation of material reason relations that can be captured by logical vocabulary can be semantically interpreted in implication-space terms. This novel correlation of proof-theoretic and model-theoretic specifications of reason relations articulates a more fine-grained self-consciousness of the semantogenic rational relations that confer conceptual content on the sentences that stand in those relations.

Pufendorf Lectures 2025: **Vocabularies of Reason**

Handout for Lecture III:

Roles and Reasons

Bob Brandom

A *vocabulary* is a pair <L,**I**> of a *lexicon* L, which is a set of sentences, and distinguished set **I** of pairs of sets of sentences of L. <X,Y>∈**I** means that the implication with premises X and conclusion Y is a good one. Premise-sets are read *con*junctively, and conclusion-sets are read *dis*junctively. We mark the *incoherence* of the premise-set X by <X,∅>∈**I**.

The *implication space* defined by a vocabulary with lexicon L is the set $\mathcal{P}$(L)x$\mathcal{P}$(L) of pairs of sets of sentences of L.
The points of the implication space are thought of as *candidate implications.* The good implications, according to the implication-space frame defined from a vocabulary, are just those in **I**.

"Kant was on the right track when he insisted that just as concepts are essentially (and not accidentally) items which can occur in judgments, so judgments (and, therefore, indirectly concepts) are essentially (and not accidentally) items which can occur in reasonings or arguments."
["Inference and Meaning" [I-4], in Kevin Scharp and Robert Brandom (eds.) *In the Space of Reasons*: *Selected Essays of Wilfrid Sellars* [Harvard University Press, 2007].

- The *extension* of a candidate implication <X,Y> is just its *goodness* value: whether or not it is a good implication (to be found in **I**).
- The *intension* or semantic interpretant of a candidate implication is its *range of subjunctive robustness* (RSR): the set of additions to its premise- and conclusion-set that *complete* it, in the sense that they would *make* it good if it is *not* good, and *keep* it good, if it *is* good.

  RSR<Γ,Δ> =$_{df.}$ {<X,Y>∈LxL : <Γ∪X,Δ∪Y>∈$\mathbf{I}_M$.
- The *implicational role* of implication Γ|~Δ is the equivalence class of (sets of) candidate implications that share its range of subjunctive robustness (intension):

  $\mathcal{R}$({<Γ,Δ>}) =df. {<X,Y>∈LxL : RSR<X,Y> = RSR<Γ,Δ>}.
- The *conceptual* (propositional) *content* of sentence A∈L is the pair of the implicational roles of its *premissory* and *conclusory* seed implications <A,∅> and <∅,A>:

  [A] = <a^+,a^-> = < $\mathcal{R}$({<A,∅ >}), $\mathcal{R}$({<∅,A>}) >.

These are the inferential *consequences* and *circumstances* of application of the sentence A.

Operations on Implicational Roles:

Symjunction: $\mathcal{R}(X) \sqcap \mathcal{R}(Y) =_{df.} \mathcal{R}(X \cup Y)$.

Adjunction: $\mathcal{R}(X) \sqcup \mathcal{R}(Y) =_{df.} \mathcal{R}(\{\Gamma \cup \Delta: \Gamma \in X, \Delta \in Y\})$.

The implication-space semantic definitions of the connectives of NMMS:

$[A] =_{df.} \langle a^+, a^- \rangle$ $[B] =_{df.} \langle b^+, b^- \rangle$

$\sqcup$ is adjunction of implicational roles, $\sqcap$ is symjunction of implicational roles

Computing the conceptual roles of logically complex sentences:

$[\neg A] =_{df.} \langle a^-, a^+ \rangle$.

$[A \rightarrow B] =_{df.} \langle a^- \sqcap b^+ \sqcap (a^- \sqcup b^+), a^+ \sqcup b^- \rangle$.

$[A \& B] =_{df.} \langle a^+ \sqcup b^+, a^- \sqcap b^- \sqcap (a^- \sqcup b^-) \rangle$.

$[A \vee B] =_{df.} \langle a^+ \sqcap b^+ \sqcap (a^+ \sqcup b^+), a^- \sqcup b^- \rangle$.

Fact: These implication-space connective definitions in terms of symjunction and adjunction of roles produce a *sound and complete semantics* for the universally LX logic NMMS—the substructurally forgiving version of classical logic.

The Metalogical Correspondence between Implication-Space and Sequent-Calculus MVs:

1. The *first* element in the roles defined by the semantic clauses corresponds to the *left* rule in the sequent calculus, and the *second* element corresponds to the *right* rule in the sequent calculus.
2. The roles super-scripted with a "+" stem from sentences that occur on the *left* in a top sequent, and the roles super-scripted with a "−" stem from sentences that occur on the *right* in a top sequent.
3. An *adjunction* $\sqcup$ indicates that the adjoined roles stem from sentences in a *single* top sequent. And a *symjunction* $\sqcap$ indicates that the symjoined roles stem from sentences that occur in *different* top sequents.

Given that the contexts Γ,Δ are always shared in all the sequents of any rule application, using this correspondence, the implication-space semantic clauses above uniquely determine the sequent rules of any logic, and the other way around.

Connective Rules of NMMS:

$L\neg$: $\dfrac{\Gamma \mid\sim \Delta, A}{\Gamma, \neg A \mid\sim \Delta}$ $R\neg$: $\dfrac{\Gamma, A \mid\sim \Delta}{\Gamma \mid\sim \Delta, \neg A}$

$L\rightarrow$: $\dfrac{\Gamma \mid\sim \Delta, A \quad B, \Gamma \mid\sim \Delta \quad B, \Gamma \mid\sim \Delta, A}{\Gamma, A \rightarrow B \mid\sim \Delta}$ $R\rightarrow$: $\dfrac{\Gamma, A \mid\sim B, \Delta}{\Gamma \mid\sim A \rightarrow B, \Delta}$

$L\&$: $\dfrac{\Gamma, A, B \mid\sim \Delta}{\Gamma, A \& B \mid\sim \Delta}$ $R\&$: $\dfrac{\Gamma \mid\sim \Delta, A \quad \Gamma \mid\sim \Delta, B \quad \Gamma \mid\sim \Delta, A, B}{\Gamma \mid\sim \Delta, A \& B}$

L$\vee$: $\dfrac{\Gamma,A|\sim\Delta \quad \Gamma,B|\sim\Delta \quad \Gamma,A,B|\sim\Delta}{\Gamma,A\vee B|\sim\Delta}$ R$\vee$: $\dfrac{\Gamma|\sim\Delta,A,B}{\Gamma|\sim\Delta,A\vee B}$

Implicational Role Inclusions:

Premissory Role Inclusion: If $RSR(a^+)\subseteq RSR(b^+)$ then for all contexts
$\Gamma,\Delta\subseteq L[\Gamma,A|\sim\Delta \Rightarrow \Gamma,B|\sim\Delta]$, so A can be substituted for B as a premise *salva consequentia*.
Conclusory Role Inclusion: : If $RSR(b^-)\subseteq RSR(a^-)$ then for all contexts
$\Gamma,\Delta\subseteq L[\Gamma|\sim A,\Delta \Rightarrow \Gamma|\sim B,\Delta]$, so A can be substituted for B as a conclusion *salva consequentia*.
If Cut (CT) holds, then $RSR(a^+)\subseteq RSR(b^+)$ iff $RSR(b^-)\subseteq RSR(a^-)$.
'Pedro is a donkey' implies 'Pedro is a mammal,' A|~B, and
'Pedro is a donkey' can be substituted everywhere as a premise for 'Pedro is a mammal', *salva consequentia*, $RSR(a^+)\subseteq RSR(b^+)$, and
'Pedro is a mammal' can be substituted everywhere for 'Pedro is a donkey as a conclusion, *salva consequential*, $RSR(b^-)\subseteq RSR(a^-)$.
Absent that global transitivity structure, premissory and conclusory roles can diverge.

Trilogics:

The paracomplete logic K3 and the paraconsistent logic LP (Graham Priest's 'Logic of Paradox') both use the three-valued Strong Kleene connective definitions, differing only in how they define consequence.
K3 treats an implication as good iff it preserves the value 1 (true), and
LP treats it as good iff it preserves non-0 (non-false) values.
K3 invalidates Excluded Middle, and is a logic of truth-value *gaps*;
Its middle value ½ may be thought of as meaning '*neither* true *nor* false.'
LP invalidates Noncontradiction, and is a logic of truth-value *gluts*;
Its middle value ½ may be thought of as meaning '*both* true *and* false.'
It has been suggested that semantic paradoxes can be dealt with by assigning paradoxical sentences like the Liar the third truth-value of ½ and drawing consequences according to K3 (Kripke) or LP (Priest).

Fact: **K3 is the logic of *premissory* implicational role inclusions and**
LP is the logic of *conclusory* implicational role inclusions.

What shows up in the (extended) semantics of truth-values as the difference between countenancing truth-value *gaps* and *gluts* shows up in the implication-space setting as the difference between *premissory* and *conclusory* role inclusions.

Four Kinds of Rational Metavocabulary:

Extrinsic-Explanatory:

1. Bilateral Deontic Normative Pragmatic Metavocabulary.
   Key concepts: Doxastic commitments to accept/reject, expressed in speech acts of assertion/denial (generically: claimings), entitlement to which can be challenged by giving reasons *against* them (governed by incompatibilities) and defended by giving reasons *for* them (governed by implications). $\Gamma|\sim\Delta$ means one cannot be entitled to commitments to accept all of Γ and reject all of Δ.

2. Truthmaker Alethic Modal Mereological Semantic Metavocabulary.
   Key concepts: Possible/impossible states and their mereological fusions, propositions as pairs of truthmakers and falsitymakers related by Exclusion (every fusion of a truthmaker and a falsitymaker of the same proposition is an impossible state). $\Gamma|\sim\Delta$ means every fusion of a truthmaker of all of Γ with a falsitymaker of all of Δ is an impossible state.

Intrinsic-Explicative:

3. Sequent-Calculus Metavocabulary for NonMonotonic MultiSuccedent (NMMS) Logic.
   Key Concepts: Metainferential rules relating reason relations. Universal LX-ness of NMMS. It is expressively complete in that for every set of implications in any (CO-compliant) base vocabulary there is a sequent in its logical NMMS extension that is good in all and only the models elaborated from base vocabularies in which those atomic sequents hold, and *vice versa*.

4. Implication-Space Semantic Metavocabulary of Conceptual Roles.
   Key Concepts: Ranges of subjunctive robustness, implicational roles, and premissory and conclusory roles of sentences. Implicational role inclusions. General constructive correlation of operations on implicational roles and connective definitions in sequent calculi generate semantic connective rules that are sound and complete for NMMS.

We have demonstrated how to construct an isomorphism between (1) and (2), a general constructive correlation between (3) and (4) up to soundness and completeness, and how to use both (3) and (4) to capture what is common to (1) and (2).
Proposal: Define *reason relations* functionally, from above, as what can be specified in all four of these kinds of rational metavocabulary, so that they stand in these relations.

Material is from Chapter 5 of Ulf Hlobil and Robert Brandom *Reasons for Logic, Logic for Reasons: Pragmatics, Semantics, and Conceptual Roles* [Routledge, 2024].

Vocabularies of Reason
2025 Pufendorf Lectures
Lund University

Lecture III

Roles and Reasons

I. Introduction

The main aim of my first two lectures was to put on the table the concept of reason relations of implication and incompatibility, and to show how that topic appears from the points of view afforded by different ways of talking about it. From the vantage point of bilateral *pragmatic* theories of the *use* of declarative sentences to make claims and defend and challenge them by giving reasons for and against them, reason relations show up as norms determining which constellations of bilateral doxastic commitments to accept or reject a speaker can be jointly entitled to. From the vantage point of truth-maker representational *semantic* theories of the *meaning* or *content* of declarative sentences, reason relations show up as alethic modal constraints on the *compossibility* of truth-making states for premises and falsity-making states for conclusions. And from the vantage point of sequent calculi codifying metainferential rules for computing new reason relations from old ones, they show up as what is both *elaborated* from the reason relations of a base vocabulary and *made explicit* by conditional and negating logical locutions (along with the aggregative Boolean helper-monkeys of conjunction and disjunction) in a logical extension of that base vocabulary.

Throughout, I have held out the prospect of understanding the claimable conceptual contents expressed by declarative sentences in terms of the reason relations they stand in to one another. This aspiration is in service of a top-down explicative strategy. Just as reason relations are to be specified "from above," triangulating on them in the terms of pragmatic, semantic, and

logical rational metavocabularies, so those reason relations are to be called upon in turn to articulate functional definitions of *propositions* in terms of the roles those sentential conceptual contents play in implications and incompatibilities relating them. Today I want to begin to make good on that promise by cashing out the metaphor of 'roles.' Yesterday's theme was *logical expressivism.* Today's is *semantic inferentialism*.

[As with the technical results at the center of my previous two lectures, the important thing is to understand just enough of the details to appreciate the *meaning* or *significance* of the construction.]

The previous discussion gives us a place to start. Fine has a robust metaphysical conception of the propositions that stand in reason relations. They are pairs of sets of states that are eligible to serve as truth-makers and falsity-makers of sentences, in virtue of satisfying the Exclusivity condition: any mereological fusion of (exact) truth-makers with (exact) falsity-makers of the same sentence must be an impossible state. The Hlobil isomorphism at the level of reason relations between truth-maker semantics and bilateral pragmatics ensures that there is a corresponding conception of the claimables in terms of which constellations of doxastic commitments to accept and reject those claimables discursive practitioners can be jointly entitled to—basically that one cannot be entitled both to assert and to deny the same claimable. These are serviceable notions within their own semantic and pragmatic domains. But they are very *different* conceptions of propositional conceptual content. Each essentially appeals to substantive conceptions native to its setting but alien to the other: in the one case, metaphysical mereological fusion and modal possibility or impossibility of states, and in the other case deontic normative concomitance of doxastic commitments to accept or reject, and preclusions of entitlement to those commitments. These two accounts of reason relations, and so of propositional conceptual content, are shot through with the substantive semantic and pragmatic concepts they use to explain their parochial senses of 'consequence' and 'incompatibility.' What we are after now is something more abstract, something these conceptions have in common, just in virtue of the propositions each account in its own way understands as being related to one another by reason relations that are intelligible as isomorphic across the two settings.

The way to fill in that conception is to define a conception of *rational* proposition or *conceptual* propositional content entirely in terms of what I earlier called 'vocabularies'. These are abstract relational structures, consisting of a lexicon and a set of reason relations defined on that lexicon. The domain is a set of sentences (or other "bearers"). Taking our cue from the multisuccedent sequent-calculus idiom for specifying reason relations, we can represent the reason relations in a vocabulary by a set **I** of pairs of sets of lexical items, where for the pair of sets of sentences X and Y, $\langle X,Y\rangle \in \mathbf{I}$ means that the implication with premises X and conclusion Y is a good one. Understanding the premise-sets and conclusion-sets on Gentzen's model rather than Tarski's means two things. First, while in both cases premise-set are read *con*junctively, on Gentzen's model conclusion-sets are read *dis*junctively, rather than *con*junctively, as Tarski's conception of consequence does. Second, if we promise not to use this notational convenience to push incompatibility into the shadows as some kind of merely second-class reason relation, we can help ourselves to Gentzen's trick for encoding incompatibilities in the form of implications, by using empty right-hand sides. So, if the conclusion-set of a good implication is empty, that is to be read as marking the incoherence of the premise-set—and so the incompatibility of any two subsets whose union is the whole premise-set.[33] Both of these conventions turn out to have substantial technical advantages.

Vocabularies specify reason relations in the sense in which the Hlobil isomorphism shows them to be common to what shows up in deontic normative guise in bilateral pragmatics and alethic modal guise in truth-maker semantics. This isomorphism at the level of reason relations is the basis of what in the first lecture I called "bimodal conceptual realism" about reason relations: the view that when all goes well the very same reason relations that normatively govern practices of making claims and rationally challenging and defending them can be understood as modally articulating the world that is thereby talked *about*. We have already seen how to do *logic* with vocabularies in this spare, technical sense. For the structural principles and connective definitions of sequent-calculus metavocabularies operate on a base vocabulary in this sense to compute both the lexicon and the reason relations of a super-vocabulary of it. Its lexicon is a superset of the base lexicon, produced by adding logical compounds of base sentences, and its reason relations are conservatively elaborated by the sequent rules from the reason relations of the base, so forming a superset of those reason relations. The question I am addressing today is how to define a formally tractable and philosophically useful concept of

[33] The utility of this notational encoding of incompatibilities as implications does not require Explosion. If $\Gamma|\sim\varnothing$ in a monotonic setting, then Γ implies everything, since we can weaken the right-hand side with any sentence or set of sentences. Absent MO, explosion in this sense does not follow from implying the empty set.

propositional conceptual content that appeals only to this minimal foundation of the reason relations of arbitrary material (nonlogical) base vocabularies.

II. Implication-Space Semantics

The construction I will present is *implication-space semantics*. It was originally adapted by ROLE-contributor Daniel Scott Kaplan from Jean-Yves Girard's phase-space semantics for linear logic, and was then developed further by Ulf Hlobil for our book.[34] Implication-space models are just another version of vocabularies. An implication space is formed from a lexicon of sentences as the set of all the *candidate implications* on that lexicon. If the lexicon is the set of sentences L, this is the set of all *pairs* of sets of sentences of L. Each element of the implication space generated by a lexicon is thought of as the pair of the premise-set and the conclusion-set of a candidate implication. An implication-space *frame* is then an implication-space together with a distinguished subset **I** of it, interpreted as the *good* implications, the ones that really hold. Clearly this implication-space redescription just lightly repackages the same information that is already available when it is shaped as or put in the form of vocabularies. Implication spaces, whose points are candidate implications, give us a good way to visualize sets of reason relations, which are just subspaces of them.

We can get a philosophical hint as to which features of these structures it will be most revealing to associate with sentences as articulating their conceptual contents, by listening once again to my hero Wilfrid Sellars. We have already appropriated two lessons from him. In my first lecture I invoked his reading of Kant to motivate an inferentialist, top-down order of explication, with the quotation:

> Kant was on the right track when he insisted that just as concepts are essentially (and not accidentally) items which can occur in judgments, so judgments (and, therefore, indirectly concepts) are essentially (and not accidentally) items which can occur in reasonings or arguments.[35]

Adapting his language to that I have been employing, Sellars's criterion of demarcation for specifically *conceptual* contentfulness is situation in a "space of implications," which

[34] Originally presented in Daniel Kaplan, *Substructural Content*, 2022 University of Pittsburgh philosophy Ph.D. thesis (online at http://d-scholarship.pitt.edu/42065/). Further developed and presented in Chapter 5 of *Reasons for Logic, Logic for Reasons* [Routledge, 2024].

[35] "Inference and Meaning" [I-4], in Kevin Scharp and Robert Brandom (eds.) *In the Space of Reasons*: *Selected Essays of Wilfrid Sellars* [Harvard University Press, 2007].

normatively govern "language-language moves" or inferences.[36] I have followed him here in these methodological commitments. Sellars was also the source for thinking of the content-articulating implications as being what he called "material" inferential relations, such as that relating "Pittsburgh is to the West of New York," and "New York is to the East of Pittsburgh," which articulate the content of *nonlogical* or *prelogical* concepts such as East and West.

Sellars makes a further claim about the implications he understands as articulating conceptual content: they must be understood as *subjunctively robust*. This point can perhaps best be appreciated from the side of pragmatics, by considering what one must be able to *do* in order to count as grasping a nonlogical concept, such as lion. The first point is that one must be able in practice to distinguish candidate implications and incompatibilities that are materially good from those that are not—however incomplete and fallible that ability is. That is, one must be disposed for instance to treat the fact that a lion is very hungry as providing good reason to think that it will attack a nearby gazelle, rather than a large rock, and as a reason against expecting it to sleep lazily in the sun or flee for its life. One might have no dispositions corresponding to many such good implications and incompatibilities, and one might be mistaken about some of them. But if one makes *no* such discrimination, then one is not deploying a concept. One's grasp of the concept essentially involves making at least *some* rough and ready practical distinction between the materially good and bad implications and incompatibilities it is involved in as premise and as conclusion.

But more is required. One must also practically associate with each of the implications a *range of subjunctive robustness*. That is, one must have some sense of what differences would make a difference to the goodness of the implication. This means realizing that the lion's hunger would no longer provide a reason for expecting it to chase the gazelle if the lion or the gazelle had been struck by lightning, squashed flat by an elephant, or shot by a hunter, but that the goodness of the implication would *not* be affected by the position of a beetle on the branch of a distant tree or the day of the week being a Tuesday. Just as one must be able practically to sort candidate implications into good ones and bad ones, one must have some sense of which changes

[36] He uses the phrase "space of implications" at CDCM §108 [ref.], and introduces "language-language" moves as part of the theory of "pure pragmatics" of "Some Reflections on Language Games" in Scharp and Brandom, *In the Space of Reasons*: *Selected Essays of Wilfrid Sellars*, op. cit..

to premises or conclusions would *make* or *keep* them good. The implications that articulate the conceptual content of defeasible, nonlogical claims such as "the lion will chase the gazelle" are not what Abelard called consequences "*hic et nunc,*" here and now, which have no consequences at all for what would follow if things were even slightly different, if the premises were changed in any way. Rather, they are good implications as instances of a pattern, and to understand them one must have *some* grasp of that pattern of other, neighboring implications that would *also* be good if this one were.

The lesson that matters for our project is that there is an essentially *modal* element to reason relations as such. The semantic significance of any one implication depends on its neighbors—on the goodness of candidate implications with slightly variant premise-sets. This thought motivates the first of three ideas determining the steps needed to build an implication-space semantics. That idea is to pursue a top-down order of explication, in which semantically interpreting *implications* in terms of other implications comes before semantically interpreting *sentences* in terms of the implications they are involved in. The first semantic interpretant of a candidate implication is just its *range of subjunctive robustness*. This RSR is a matter of which additions of premises or conclusions to a candidate implication yield good ones. The RSR of a candidate implication consists of all its good implicational *completions*: the pairs of additional premises whose addition would *make* it good, if it is *not* good, or *keep* it good, if it *is* good. The range of subjunctive robustness determines the *intensional* element of the implicational role of a candidate implication, as its value as a *good* implication, or not—its goodness value—is the *extensional* component. (Compare: truth *conditions* and truth *values*.) When at the third stage in our construction we finally assign propositional conceptual contents to sentences, those contents need to be identified and individuated finely enough to respect and determine not only the extensional issue of which implications are materially good or bad (according to an implication-space frame), but also the intensional dimension of their complex ranges of subjunctive robustness.

In classical, topologically closed, specifically monotonic settings, the ranges of subjunctive robustness of each implication have the same form: if at a certain point, adding further premises or conclusions to a candidate implication yields a materially good one, then all

the other implications accessible from it by adding even more premises and conclusions are also good. In nonmonotonic settings, where adding premises or conclusions can turn a good implication into a bad one, ranges of subjunctive robustness of implications are much more complicated. In both settings, there is important information in the facts about what additional premises (and conclusions) it would take to *make* good implications out of a candidate that is *not* as it stands, good. We can adequately represent that complexity by associating with each implication, as its semantic interpretant, its range of subjunctive robustness, in the sense of all the other candidate implications that, when dual-unioned with the implication being interpreted, yield a good implication. What I am calling 'dual pointwise union' of candidate implications here is just unioning their premise-sets and unioning their conclusion-sets. The range of subjunctive robustness (RSR) of an implication <X,Y> is then the set of all candidate implications, good or bad, that when dual-unioned with it *keep* it good if it is good, or *make* it good if it is not good—all of this relative to an implication-space frame or vocabulary.

Formally: $\forall X,Y \subseteq L[\ RSR_M(\langle X,Y\rangle) =_{df.} \{\langle W,Z\rangle : W,Z \subseteq L \text{ and } \langle X \cup W, Y \cup Z\rangle \in \mathbf{I}_M\}\ ]$.[37]

The second big idea shaping implication-space semantics is that we want to group together lexically different expressions that play the same semantic role. If two candidate implications have the same range of subjunctive robustness, then they play the same implicational role, and are accordingly semantically equivalent. So, looking forward, if substituting one sentence for another never changes the range of subjunctive robustness of any implication they are involved in, then they are semantically equivalent sentences. The *implicational role* of an implication (or set of implications) can be represented by the equivalence class of implications that all have the same range of subjunctive robustness (RSR).[38] We accordingly move from understanding the subjects of semantic interpretation to be

[37] Notice that at this primary level of implications, intensions determine extensions at each implication frame. After all, the question of whether a candidate implication is a good one, whether it is in the distinguished set $\mathbf{I}_M$, is equivalent to the question of whether the minimal candidate implication $\langle\varnothing,\varnothing\rangle$ is in its RSR. For the pointwise union of that minimal candidate implication with any other candidate implication is just that implication itself. And $\langle\varnothing,\varnothing\rangle$ is in the RSR of an implication just in case dual unioning it with that implication yields a good one.

[38] For simplicity, I will talk as though implicational roles are sets implications (all those that share the same range of subjunctive robustness). In fact for technical reasons we need to define roles also for *sets* of implications, where the RSR of a set of implications is the intersection of the RSRs of its elements. Sets of implications can share their RSR with single implications, so implicational roles are really *sets of sets* of implications. This complication matters for defining the operations on roles: adjunction and symjunction. I suppress these details, in my story here. Each of my not-quite-correct formulations can be replaced by a more cumbersome accurate one.

implications in the sense of pairs of sets of sentences of a lexicon to understanding the subjects of semantic interpretation to be implicational roles of implications, which are understood as equivalence classes of such implications assimilated accordingly as they have the same range of subjunctive robustness.

Implicational roles are very special equivalence classes of implications. Not every set of implications *is* an implicational role. But every set of implications *has* an implicational role: the equivalence class of all the (sets of) implications that have the same range of subjunctive robustness as that set. For all sets of implications have ranges of subjunctive robustness (the intersection of the ranges of the individual RSRs), and roles are just sets of (all the sets of) implications that have that RSR. These implicational-role equivalence classes of implications are the building-blocks of implication-space semantics. Henceforth our principal concern is with these implicational roles. We will assemble the propositional conceptual contents expressed by sentences out of implicational roles, and define and exploit operations for forming new implicational roles from old ones.

This shift in conceptions of what is semantically interpreted crucially lifts the semantic discourse to a higher level of abstraction. When we considered Fine's truth-maker semantics in my first lecture, we saw that worldly propositions in his sense—pairs of sets of truth-making and false-making mereological states that meet his Exclusivity condition—can stand to one another in reason relations of implication and incompatibility. Treating a universe of propositions in this sense as the lexicon of a vocabulary, we can generate implication-space frames from those vocabularies, and compute the ranges of subjunctive robustness and (so) conceptual roles of truth-maker implications. The abstraction achieved by treating implications as equivalent if they have the same ranges of subjunctive robustness produces a notion of <u>implicational role</u> that applies equally well to the truth-maker setting. That means that its implications, and eventually, its worldly propositions, can be understood as playing the *very same* implicational conceptual roles as those played by the sentences of a linguistic vocabulary as used to make claims and give reasons for and against them. Implicational roles capture what is *common* to the truthmaker alethic modal mereological semantics and to the bilateral deontic normative pragmatics I talked about in my first lecture.

The third stage in articulating implication-space semantics is the extension of the concepts of range of subjunctive robustness and so implicational role from semantically interpreting candidate *implications* (and sets of them) to semantically interpreting *sentences* (and other bearers of conceptual content, such as truth-maker propositions). This move from interpreting implications to interpreting sentences is of the essence of the top-down order of semantic explication that I have been pursuing from the beginning. To implement that strategy, we must understand sentences semantically in terms of sets of implications: specifically, in terms of the equivalence classes of implications that are implicational roles. Here the key thought is that each sentence in a vocabulary is most naturally associated with *two* implicational roles: the roles of the good implications in which it appears as a *premise*, and the roles of the good implications in which it appears as a *conclusion*. This conception is a descendant of Dummett's way of thinking about propositional contents in terms of the pair of a sentence's appropriate *consequences* of application and its *circumstances* of appropriate application(which I adapted and developed in *Making It Explicit*).[39]

The premissory and conclusory roles are *different* sets of implications.[40] One is determined by the good implications in which the index sentence shows up as a premise and the other is determined by the good implications in which it appears as a conclusion. The implication-space apparatus of ranges of subjunctive robustness provides a simple way to represent those roles. The set of good implications in which the sentence A appears as a premise is interdefinable with RSR<A,∅>. For by definition <X,Y>∈RSR<A,∅> just in case adding A to X yields a good implication: <X∪{A},Y>∈$\mathbf{I}_M$. And dually, the set of good implications in which the sentence A appears as a conclusion is determined by RSR<∅,A>. We can call <A,∅> and <∅,A> the premissory and conclusory *seed* implications of A, and RSR(<A,∅>) and RSR(<∅,A>) the *premissory* and *conclusory* RSR-sets of implications of A.

[39] The inferential *circumstances* of appropriate application of a sentence are represented by its *conclusory* role, and the appropriate inferential *consequences* of application of a sentence are represented by is *premissory* role.

[40] The overlap between these two sets consists entirely of good implications in any frame that satisfies the minimal structural condition of Containment, in which some conclusion appears also as a premise—which most of those we are concerned with do.

The role $\mathcal{R}(\{<A,\varnothing>\})$ is the set of all implications that are intersubstitutable *salva consequentia* with A as premise, and similarly, $\mathcal{R}(\{<\varnothing,A>\})$ is the set of all implications that are intersubstitutable *salva consequentia* with A as a conclusion. But what does it mean for an *implication* to be intersubstitutable with a *sentence*? Associating each sentence in the lexicon with native sets of implications is the trick that makes possible a top-down order of semantic articulation: semantically identifying sentences with sets of implications—in our case, with *pairs* of sets of implications. What intuitive sense can we make of the semantic equivalence of sentences with sets of implications (the ranges of subjunctive robustness of their seeds)? The thing to focus on is the relations between implications that play the same implicational role, in that they have the same range of subjunctive robustness.

Suppose $<\Gamma,\Delta>$ is a candidate implication that plays the same implicational role as —and so has the same RSR as—the premissory seed $<A,\varnothing>$ of the sentence A. So any candidate implication $<X,Y>$ is in the RSR of $<\Gamma,\Delta>$ just in case it is in the RSR of $<A,\varnothing>$. But that means that adding A as a premise to X yields a good implication $<X\cup\{A\},Y>$ just in case adding Γ to the premise *and* adding Δ to the conclusions to the same context *also* yields a good implication, $<X\cup\Gamma,Y\cup\Delta>$. To get the effect of weakening the implication by adding A to the premises, one must weaken *both* sides, adding Γ to the premises *and* Δ to the conclusion. In this sense, the candidate implication $<\Gamma,\Delta>$ is intersubstitutable with A *as a premise*, *salva consequentia*. This candidate *implication*, $<\Gamma,\Delta>$[41] means the same as, plays the same role in implications as the *sentence* A does, *as a premise.*

This is the explanatory route from the top down: from ranges of subjunctive robustness and (so) implicational roles of implications down to RSRs and conceptual roles of sentences, with the complication that sentences correspond to *pairs* of (sets of) implications. In this way we vindicate the distinctive inferentialist conviction that *sentences* should be understood semantically in terms of their role in *implications*—indeed, their inferential circumstances and consequences of application. Those roles are defined in terms of the distinctive kind of *modality*

[41] Which implication might or might not be a good one, relative to a frame, just as we can assign truth conditions to sentences in standard semantics without asking whether or not they are true in some particular model.

that is built into reason relations, in the form of *ranges of subjunctive robustness*.[42] Both the premissory and the conclusory roles of any sentence are classes of implications that are role-equivalent, in the sense of intersubstitutable *salva consequentia* (as codified in ranges of subjunctive robustness).[43]

In this way every sentence A of the lexicon from which the implication-space of a frame derives is associated with *two* implicational roles, two equivalence classes of implications: a premissory role consisting of all the implications that share a range of subjunctive robustness with the seed candidate implication that has the singleton {A} as its premise-set and the empty set as its conclusion set, and a conclusory role consisting of all the implications that share a range of subjunctive robustness with the seed candidate implication that has the empty set as its premise-set and singleton {A} as its conclusion set. But these are far from all the propositional contents that can be constructed from the implicational roles of a particular implication frame. For we can ask: what constraints are there on the choice, from the set of implicational roles of an implication frame, of two of them to serve as the premissory and conclusory roles of a propositional content? Is there some condition each must meet, other than just being a conceptual role, in order to function as premissory or conclusory roles of some propositional content? Must the two conceptual roles exhibit some sort of Dummettian "harmony" in order to function as a well-behaved propositional content?

Well, what do we expect of a "well-behaved" propositional content? The idea we have been working with from the beginning is that propositional contents should be understood to be what stand to one another in reason relations of implication and incompatibility. Those, we think, can be explained in *both* bilateral pragmatic and truth-maker semantic metavocabularies. Further, in the light of the logical expressivism I elaborated last time, we might want to add, as a sort of acid test of standing in such reason relations, that propositional contents can occur embedded as the antecedents of conditionals that make explicit those implication relations and

[42] From the perspective of this order of explication, this is the fundamental kind of modality: conceptual, in the sense of pertaining to reason relations, which the Hlobil isomorphism shows can be understood as common to or amphibious between what is expressed by alethic and deontic locutions.

[43] Put more carefully, they are sets of sets of implications, whose singleton elements are intersubstitutable with A as premise or as conclusion.

negated conditionals that make explicit the incompatibility relations. It turns out that we can satisfy all those criteria of adequacy for propositional contents formed by pairing *arbitrary* implicational roles as premissory and conclusory roles of propositional contents. We need put no restrictions at all on the choice of implicational roles to serve as premissory and conclusory roles of a proposition, not even relational ones codifying a notion of harmony between them. The full set of propositional contents definable on an implication frame can safely be taken to be the whole set of all ordered pairs of implicational roles of that frame. In fact, the cases we care about satisfy the very weak Containment condition, which ensures that every sentence implies itself. It follows from Hlobil's isomorphism that this minimal, but still substantive, constraint corresponds directly to Fine's Exclusivity condition relating truthmakers and falsitymakers of the same sentence.

III. Operating on Implicational Roles

We are in a position to appreciate a remarkable but subtle consequence of considering the conceptual contents determined by *all* the pairs of implicational roles. That is that the set of good implications $\mathbf{I}_M$ of an implication frame not only determines the propositional contents expressed by sentences of the lexicon that generated the space of candidate implications, but also *many* more propositional contents that are *not* expressed by any sentences of the lexicon, but whose roles are determined entirely by the reason relations among sentences that *are* in the lexicon on which both the vocabulary and the implication-space are based.[44] The reason relations among sentences of the lexicon generate a substantial semantic surplus of propositional contents determined by those reason relations, but not expressible within the lexicon. This hidden semantic territory opened up to view by the original reason relations among sentences of the lexicon is a complex landscape. Some of its denizens are massively defective, combining, say, the premissory role of 'That is a donkey,' with the conclusory role of 'the pond is icing over,' or 'electrons are not composed of quarks.' They either fail to imply themselves or underwrite bad implications, which makes such contents unlikely to be of much epistemic use. But other bits of this semantic shadow matter, other unexpressed propositional contents, are of the utmost importance. It is here, for instance, that the propositional contents of logical compounds of the contents expressed by sentences of the lexicon are to be found.

The point I want to emphasize is that arbitrarily constructed propositional contents that are not expressed by sentences of the lexicon are wholly determined by the reason relations of the implication frame they live in, and only the sentences of the lexicon of that frame appear in their ranges of subjunctive robustness and implicational roles. Every implication frame generates such a shadow realm of expressible-but-unexpressed propositional contents. Sentences expressing *these* propositions can be added without changing *any* of the reason relations of the vocabulary. What makes something a specifically *propositional* content is that it stands in

[44] Fine's worldly propositions also in general outrun what is expressible in any standard kind of language. But that is because of the richness of his mereological modal metaphysics of states. What is special to the implication-space setting is that the semantic interpretants themselves are constructed entirely from the reason relations among the sentences of a base vocabulary, yet include many propositions not expressed by those sentences.

reason relations to other such contents. Since shadow propositions are not expressed by sentences of the lexicon, they do not occur in any good implications in the set $\mathbf{I}_M$ of the frame they inhabit. But their premissory and conclusory roles consist entirely of sentences of the lexicon and we can define entailments (and incompatibilities) just in these terms. For we can understand a set of premises Γ semantically entailing a set of conclusions Δ just in case the empty sequent $\langle\varnothing,\varnothing\rangle$ is in the range of subjunctive robustness of $\langle\Gamma,\Delta\rangle$.[45]

An important metasemantic criterion of adequacy of an assignment of semantic interpretants in a formal semantics is that there should be a uniform procedure for computing the semantic interpretants of lexically complex sentences from the semantic interpretants of lexically simple ones—so enabling the derivation of the reason relations governing the use of the sentences of the more complex vocabulary from those of the base. This is where atomistic, bottom-up orders of semantic explication shine. Being able to account for the meanings of more complex sentences in terms of the meanings of simpler ones is the main strength and principal *raison d'être* for this sort of approach. By contrast, the criterion of adequacy that a formal semantics must recursively determine semantic interpretants for complex sentences from those of simpler sentences is potentially much harder to satisfy in the context of a top-down order of explication. The thoroughgoing *holism* of our approach, intensified by the aspiration to handle even radically open-structured reason relations, at least makes more difficult the kind of recursive determination of the semantic interpretants of more complex sentences from simpler ones, which is the strength of bottom-up atomistic approaches. Jerry Fodor took it to be impossible, and that that impossibility showed the bankruptcy of holistic semantic approaches, including especially inferentialist ones.[46]

The operators for forming complex sentences from simpler ones that we understand best are *logical* operators. That is why the standard initial test-bench for a formal semantics (for instance, for Frege, Tarski, Kripke, and Fine) is its application to specifically *logical* vocabulary. We

---

[45] Here it matters that implicational roles are officially sets of *sets* of implications. This condition is equivalent to "every element in every set of the implicational role of $\langle\Gamma,\Delta\rangle$ is in $\mathbf{I}_M$," and "the union of every set in the implicational role of $\langle\Gamma,\Delta\rangle$ is in $\mathbf{I}_M$." A corresponding definition applies to semantic incompatibility.

[46] Jerry Fodor, Ernest LePore *Holism, a Shopper's Guide* [Wiley-Blackwell, 1992], and Jerry Fodor, Ernest LePore "Brandom's Burdens: Compositionality and Inferentialism" *Philosophy and Phenomenological Research* Vol. LXIII, No. 2, September 2001, pp. 465-491.

know how to use metainferential rules couched in a sequent-calculus metavocabulary to compute the reason relations of logically extended vocabularies from the reason relations of the base vocabularies to which logical sentential operators are added. The question is then whether there are uniform procedures for deriving the premissory and conclusory roles of logically complex sentences from those of their component sentences, so that those roles determine the right reason relations among logically complex sentences: the very same reason relations computed in quite a different way by sequent-calculus rules. This is a very definite question, with clear technical criteria of adequacy. I am happy to be able to report that the answer is 'Yes.'

The implication-space interpretation of a sentence A can be unpacked like this.
We symbolize the implicational role of A by enclosing A in square brackets:
[A].
We can decompose that into the pair of a premissory implicational role and a conclusory implicational role:
[A] = $\langle a^+, a^- \rangle$.
Each of those elements can be further decomposed:
[A] = $\langle a^+, a^- \rangle$ = $\langle$ ℛ ({$\langle A, \varnothing \rangle$}), ℛ ({$\langle \varnothing, A \rangle$}) $\rangle$.

We want to look at what natural operations there are combining the implicational roles we are treating as semantic interpretants of sentences. Because we keep separate sets of books on the premissory and conclusory implicational roles, one natural operation is to create a sort of converse implicational role by swapping these.

The negation rules of our favored expressivist logic NMMS are just those of standard classical logic (Gentzen's LK):

$$L\neg: \quad \frac{\Gamma \mid\sim \Delta, A}{\Gamma, \neg A \mid\sim \Delta} \qquad\qquad R\neg: \quad \frac{\Gamma, A \mid\sim \Delta}{\Gamma \mid\sim \Delta, \neg A}$$

We can read the left rule as saying that the role of ¬A as premise is the same as the role of A as conclusion, and the right rule as saying that the role of ¬A as conclusion is the same as the role

of A as premise. That is to say that $[\neg A] = \langle a^-, a^+ \rangle$. For, more specifically, the left rule says that any candidate implication $\langle \Gamma, \Delta \rangle$ that yields a good implication when A is added as a conclusion—that is, anything in the conclusory implicational range of A—will yield a good implication when $\neg A$ is added as a premise—that, is it will be in the premissory implicational range of $\neg A$. The metainferential rules that relate the occurrence of A on one side of the turnstile to $\neg A$ on the other are equivalent to defining the implicational role of $\neg A$ as the converse, in this sense, of the implicational role of A.

To go further, in addition to this one-place operation of swapping premissory and conclusory implicational roles to produce a kind of semantic converse or inverse corresponding to negations codifying incompatibility, we need two-place operations that *combine* different implicational roles to make new ones that are compounds of the originals, in order to interpret conditionals, conjunctions, and disjunctions. The Boolean algebras that interpret classical logical connectives in topologically closed settings appeal at this point to operations of *unioning* and *intersecting* the sets that are assigned to sentences as their semantic interpretants—whether those propositional contents are understood as sets of models, or possible worlds, or truth conditions, or whatever. We can adopt this general idea, while acknowledging that adapting it to work in radically open-structured settings will require some adjustments.

The example of negation shows that the premissory role of a logical compound can depend on the conclusory role of one of its components, and *vice versa*. So we should think of these operations as applying to a pair (more generally, a set) of implicational roles, whether premissory or conclusory, and determining another implicational role, whether premissory or conclusory. The happy complication that makes our construction possible is that because we start our semantic interpretation with (candidate) implications rather than sentences, there are actually two loci in our semantic interpretants to which union-like and intersection-like operations could be applied.

We have two basic operations on roles of sets of implications, corresponding roughly to intersection and union operations on sets:

Symjunction: $\mathcal{R}(X) \sqcap \mathcal{R}(Y) =_{df.} \mathcal{R}(X \cup Y)$.

Adjunction: $\mathcal{R}(X) \sqcup \mathcal{R}(Y) =_{df.} \mathcal{R}(\{\Gamma \cup \Delta : \Gamma \in X, \Delta \in Y\})$.

The symjunction of two roles is just the role of their union, and the adjunction is the role of the set consisting of all the pointwise unions of elements of the one set with elements of the other. Although both role operations are defined using set-theoretic union, the effect at the level of ranges of subjunctive robustness is that of intersection for symjunction, since the RSR of a set of implications is defined as the intersection of the RSRs of its elements.

With these operations on board, we can formulate semantic definitions of sentential logical connectives, by showing how to compute the implicational roles of logically complex sentences from the implicational roles of their component sentences.

Here is one set of such implication-space semantic definitions of logical connectives:

$[A] \quad =_{df.} \quad \langle a^+, a^- \rangle \qquad [B] \quad =_{df.} \quad \langle b^+, b^- \rangle$

$\sqcup$ is adjunction of implicational roles, $\sqcap$ is symjunction of implicational roles

$[\neg A] \quad =_{df.} \quad \langle a^-, a^+ \rangle.$

$[A \rightarrow B] \quad =_{df.} \quad \langle a^- \sqcap b^+ \sqcap (a^- \sqcup b^+), \ a^+ \sqcup b^- \rangle.$

$[A \& B] \quad =_{df.} \quad \langle a^+ \sqcup b^+, \ a^- \sqcap b^- \sqcap (a^- \sqcup b^-) \rangle.$

$[A \vee B] \quad =_{df.} \quad \langle a^+ \sqcap b^+ \sqcap (a^+ \sqcup b^+), \ a^- \sqcup b^- \rangle.$

The right-hand side of each of these definitions specifies an ordered pair of RSR-equivalence classes of implications, which are the premissory and conclusory implicational roles of the logically compound sentences. And it does so entirely in terms of operations of adjunction and symjunction applied to the implicational roles (premissory and conclusory) of their component sentences.

It will perhaps come as a relief to hear that I am not going to try to motivate these definitions in detail. I present the definitions so that you can see a bit of implication-space semantics in action. But the philosophically important point is this: **These implication-space semantic definitions of the connectives provide a sound and complete semantics**, in a very strong sense, **for the universally LX logic NMMS** that I presented last time. Not only do they determine the same set of purely logical reason relations, but given *any* nonlogical base vocabulary, these semantic rules determine exactly the same reason relations of implication and incompatibility for the logical extension of that vocabulary as those that are determined by the

metainferential sequent-calculus rules we used to introduce the logic NMMS. These consequences include all those implications and incompatibilities among logically complex sentences that do *not* hold in virtue of logic alone, but depend on, reflect, and express the idiosyncrasies of the material reason relations of the underlying base vocabulary.

The *semantic completeness* of implication-space semantics for the logic NMMS means that there are two ways to compute the reason relations among sets of logically complex sentences from the reason relations of a logically atomic base vocabulary. One can do so directly, using the connective definitions of NMMS specified in a sequent-calculus metavocabulary. Or one can first compute ranges of subjunctive robustness and implicational roles for sets of implications in the base vocabulary, use those to define premissory and conclusory roles for atomic sentences, combine those roles according to the *semantic* clauses for NMMS connectives, using role operations of adjunction, symjunction, and inversion, and then determine reason relations from those roles for logically complex sentences by the definition of semantic entailment in terms of roles. The semantic completeness result guarantees that the reason relations that result from these two procedures exactly coincide, for every base vocabulary. (The diagram commutes.) That is what I meant by referring to the *strength* of the correspondence between the functions that compute one set of reason relations from another specified in our proof-theoretic sequent-calculus metavocabulary and the functions on roles that compute one set of reason relations from another specified in our model-theoretic implication-space semantic metavocabulary.

We saw last time that NMMS is universally LX. The *expressive completeness* of NMMS means that, computed in either the sequent-calculus way or the implication-space way in terms of conceptual roles, the logically complex sentences *completely* codify the reason relations, of the base vocabulary *and* the logically extended vocabulary, in the strong sense that for *every* set of base sequents, there is a sequent in the logically extended vocabulary that holds just in case those sequents hold in the base, and *vice versa*. In this sense, the reason relations of any base vocabulary can be made fully explicit in the form of logically complex sentences of the extended vocabulary. With their reason relations computed either way, conditionals still express implications and, together with negation, incompatibilities.

And, just as in the sequent-calculus case NMMS works smoothly to elaborate and explicate even open-structured, nonmonotonic and nontransitive reason relations, so too the semantic definitions of the implicational roles of logically complex sentences in terms of adjunctions and symjunctions of the implicational roles of their components elaborate and explicate radically substructural implication-space models. Open-structured reason relations, which are not universally monotonic or transitive—and even hypernonmonotonic consequence relations with implications that cannot even in general be weakened by adding their own consequences to their premise-sets—are incorporated into the conceptual roles of sentences of the logically atomic base vocabularies in way that is faithfully reflected in the conceptual roles of the logically complex sentences that make those consequences explicit, which are derived from them. The conceptual roles conferred even by radically substructural reason relations still combine with one another logically in just the ways required for expressive completeness, both of the underlying material reason relations and of the reason relations among logically complex sentences that derive from them.

It is worth comparing these semantic definitions of the roles of logically complex sentences in terms of adjunctions and symjunctions of the roles of their components with the sequent-calculus definitions of those connectives:

Connective Rules of NMMS:

$L\neg$: $\dfrac{\Gamma \mathrel{|\sim} \Delta, A}{\Gamma, \neg A \mathrel{|\sim} \Delta}$ $\quad$ $R\neg$: $\dfrac{\Gamma, A \mathrel{|\sim} \Delta}{\Gamma \mathrel{|\sim} \Delta, \neg A}$

$L\rightarrow$: $\dfrac{\Gamma \mathrel{|\sim} \Delta, A \quad B, \Gamma \mathrel{|\sim} \Delta \quad B, \Gamma \mathrel{|\sim} \Delta, A}{\Gamma, A \rightarrow B \mathrel{|\sim} \Delta}$ $\quad$ $R\rightarrow$: $\dfrac{\Gamma, A \mathrel{|\sim} B, \Delta}{\Gamma \mathrel{|\sim} A \rightarrow B, \Delta}$

$L\&$: $\dfrac{\Gamma, A, B \mathrel{|\sim} \Delta}{\Gamma, A \& B \mathrel{|\sim} \Delta}$ $\quad$ $R\&$: $\dfrac{\Gamma \mathrel{|\sim} \Delta, A \quad \Gamma \mathrel{|\sim} \Delta, B \quad \Gamma \mathrel{|\sim} \Delta, A, B}{\Gamma \mathrel{|\sim} \Delta, A \& B}$

$L\vee$: $\dfrac{\Gamma,A \mid\sim \Delta \quad \Gamma,B \mid\sim \Delta \quad \Gamma,A,B \mid\sim \Delta}{\Gamma,A\vee B \mid\sim \Delta}$ $R\vee$: $\dfrac{\Gamma \mid\sim \Delta,A,B}{\Gamma \mid\sim \Delta,A\vee B}$

There is a robust *correlation* between semantic clauses computing implicational roles from implicational roles and sequent-calculus meta-inferential rules, which extends far beyond the soundness and completeness of NMMS:

- The *first* element in the roles defined by the semantic clauses corresponds to the *left* rule in the sequent calculus, and the *second* element corresponds to the *right* rule in the sequent calculus.
- The roles super-scripted with a "+" stem from sentences that occur on the *left* in a top sequent, and the roles super-scripted with a "−" stem from sentences that occur on the *right* in a top sequent.
- An *adjunction* indicates that the adjoined roles stem from sentences in a *single* top sequent. And a *symjunction* indicates that the symjoined roles stem from sentences that occur in *different* top sequents.

Given that the contexts are always shared in all the sequents of any rule application, using this correspondence, the semantic clauses above uniquely determine the sequent rules of NMMS, and the other way around.

As an example, consider the promissory role of $A\rightarrow B$, which is

$[A\rightarrow B] \quad =_{df.} \quad < a^{-}\sqcap b^{+}\sqcap(a^{-}\sqcup b^{+}), ...>$

And the corresponding left rule of the sequent calculus.

$L\rightarrow$: $\dfrac{\Gamma\mid\sim\Delta,A \quad B,\Gamma\mid\Delta \quad B,\Gamma\mid\sim\Delta,A}{\Gamma,A\rightarrow B\mid\sim\Delta \; \Gamma,A\mid\sim B,\Delta}$

Only the top line of the sequent-calculus rule matters here, since it gives the premises of the metainference. The implication-space formulation has three parts, joined by symjunctions, corresponding to the three premises of $L\rightarrow$. The first 'A' appears on the right-hand side of its sequent, so gets a minus. The 'B' in the second sequent is on the premise-side of the turnstile, so it gets a plus. The third sequent has both a 'B' on the premise side and an 'A' on the conclusion side, so the 'a' gets a minus, the 'b' a plus, and those roles are adjoined.

The important point is that the correspondence is not specific to NMMS. These principles show how to translate back and forth between essentially *any* sequent-calculus specification of the reason relations of logically complex sentences and an implication-space semantic specification of those same reason relations. This is how to compute implication-space semantic operations on roles for logics whose rules are specified in sequent-calculus metavocabularies—and *vice versa*. This metalogical correspondence between two vocabularies for specifying reason relations should be understood as a crucial datapoint in the ongoing attempt to understand the relative expressive powers of proof-theoretic and model-theoretic rational metavocabularies.

There is also a straightforward way to construct from any truth-maker model an implication-space frame that validates exactly the same consequence and incompatibility relations, and *vice versa*. So any logic specifiable in truth-maker terms is specifiable in implication-space terms.[47] In particular, the expressive power of the truth-maker setting to codify *hyperintensional* logical and semantic relations is also reproducible with implication frames. It is further possible to extend the implication-space framework by moving from sets to multisets to handle *noncontractive* logics, and generalizing the correlation to sequent rules whose contexts are not shared, which allows the treatment in implication-space semantics of multiplicative and additive linear logics (MALL).[48]

[47] *RLLR* 5.2.2, pp. 224 ff.
[48] *RLLR* 5.4.2 (pp. 232-237) and *RLLR* Section 5.7.5, pp. 263 ff..

IV. Reconstructing Three-Valued Logics

Implicational roles are equivalence classes of implications that share a range of subjunctive robustness.[49] When two implications share their range of subjunctive robustness, and hence play the same implicational role, that means that they are intersubstitutable everywhere *salva consequentia.* Substituting one for the other never turns a good implication into a bad one. It can also happen that the range of subjunctive robustness of one implication (or set of them) is a proper subset of the range of subjunctive robustness of another. Then the first can be substituted for the second *salva consequentia*, but not the other way around. We call these asymmetric relations "role inclusions" (even though the inclusion relations are really among the RSRs that define the roles, not the roles themselves). When we consider role inclusions for propositional contents, we have to consider both the role-inclusions of the proposition's premissory roles, and the role-inclusions of its conclusory roles.

In a structurally closed vocabulary, if A implies B, then it is also true that substituting A for B as a premise never turns a good implication into a bad one, and substituting B for A as a conclusion never turns a good implication into a bad one. Because 'Pedro is a donkey' implies 'Pedro is a mammal,' 'Pedro is a donkey can be substituted everywhere for 'Pedro is a mammal' *as a premise*, without turning any good implication into a bad one. And 'Pedro is a mammal' can be substituted everywhere for 'Pedro is a donkey' *as a conclusion*, without turning any good implication into a bad one. In a substructural vocabulary, in particular, one where transitivity fails, these notions can come apart. Then we need to keep separate track of premissory and conclusory role inclusions.[50]

---

[49] Strictly, they are sets of *sets* of implications that share ranges of subjunctive robustness, but I am suppressing this level of detail here.

[50] If A|~B then A implies B in the *internal* consequence relation. If for all contexts Γ, Δ, if Γ,B|~Δ then Γ,A|~Δ, then A implies B in the *premissory external* consequence relation (role inclusion). This means that A can replace B as a premise, saving the goodness of implications, which is tracked by K3. If for all contexts Γ, Δ, if Γ|~A,Δ then Γ|~B,Δ, then A implies B in the *conclusory external* consequence relation (role inclusion). This means that B can replace A as a conclusion, saving the goodness of implications, which is tracked by LP.

One interesting perspective afforded by the vantage point of role inclusions in implication-space semantics concerns the familiar trilogics, Graham Priest's "logic of paradox" LP and the logic K3 that Kripke appeals to in his approach to semantic paradoxes, as well as the Strict/Tolerant logic ST and its converse TS. In the semantic metavocabulary of three-truth-valued logics, these all share the Strong Kleene connective definitions, differing only in how consequence is defined. In these terms, LP understands the third truth-value as meaning 'both truth and false' and K3 understands that third value as meaning 'neither true nor false'. LP accordingly shows up as the logic of truth-value 'gluts,' which denies the universal validity of the principle of Noncontradiction, and K3 as the logic of truth-value 'gaps', which denies the universal validity of the principle of Excluded Middle.

The fact of interest here is that in the implication-space setting, **LP is just the logic of conclusory role inclusions, and K3 is the logic of premissory role inclusions**. That is, LP tells us which conclusions can always be replaced by which others, *salva consequentia*, and K3 tells us which premises can be. They are both *logics*, and so address only the consequences that hold in all suitable implication-space models. Understanding them as the logics of conclusory and premissory role inclusions reveals how they can naturally be extended to material consequence relations on *non*logical sentences.[51]

We can exploit the connections with the truthmaker setting to show also that K3 is the unilateral external "logic of verifiers," in the sense that K3 preserves compatibility with the verifiers of the premises (jointly) to the verifiers of the conclusions (separately). And LP is the unilateral "logic of falsifiers," in the sense that LP preserves the compatibility potential of the falsifiers of the conclusions (jointly) to the falsifiers of the premises (separately). So the isomorphism between the reason relations specified by the truthmaker semantics and those specified by the implicational phase-space semantics goes beyond the internal (bilateral) consequence relations all the way to the external (unilateral) consequence relations as well.

[51] Strictly, this account of LP and K3 holds only for *conic* implication-space models—a condition with close relations to monotonicity. Obviously what holds in all models holds in all conic models. And what fails in classical logic has a conic countermodel. In this sense, the conic models can be used to represent classical logic (namely as ST), which is K3 on the left and LP on the right.

What last time I called the "basic discursive bipolarity" can show up both in the form of the opposition of truth-values true/false (extensionally, and so truth-makers and falsity-makers of sentences intensionally) and in the opposition between the premises and the conclusions of implications (extensionally, and so premissory and conclusory roles of sentences intensionally). The fact that what shows up as truth-value gaps and gluts in the first sort of metavocabulary shows up as premissory and conclusory role inclusions in the second sort of metavocabulary is another crucial data point in understanding the relations between these frameworks.

V. Conclusion

In my remarks this time I have sketched the outlines of how an expressively powerful and metaconceptually enlightening implication-space semantics can be formally elaborated from the very spare representation of reason relations in a vocabulary. Just as the classical truth-based tradition envisaged assimilating sentential and other expressions accordingly as they can be intersubstituted *salva veritate* (without turning any true sentences into ones that are not true), our implication-based approach assimilates implications accordingly as they can be intersubstituted *salva consequentia* (without turning any good implications into ones that are not good). Implications and sets of implications are assigned sets of implications as semantic interpretants: implicational roles as equivalence classes of implications with the same range of subjunctive robustness. Only then, following our top-down explanatory methodology, are sentences assigned pairs of implicational roles—so, pairs of equivalence classes of implications—as the propositional conceptual contents that are their semantic interpretants. New, more abstract reason relations can then be defined by operations on implicational roles.[52]

The same union-like and intersection-like operations that produce new implicational roles from old ones then make it possible to define operations corresponding to sentential logical connectives, so as to produce a semantics that is sound and complete for the maximally expressive, universally LX, logic NMMS I introduced last time. Further, the relations between semantic operations on implicational roles and propositional conceptual contents, on the one hand, and features of sequent-calculus rules on the other turns out to be quite systematic, underwriting a general metalogical correspondence (up to soundness and completeness) of logics specified in the proof-theoretic sequent calculus metavocabulary and those same logics specified in the model-theoretic implication-space metavocabulary. The semantic side of that

[52] The idea is that a premise-set of sentences Γ semantically entails a conclusion-set Δ in case every dual fusion of an element of the premissory role of a sentence in Γ with an element of the conclusory role of a sentence of Δ is a good implication.

correspondence can be expressed even more abstractly and perspicuously in the vocabulary of implicational role inclusions.

I began the first lecture by introducing the topic of *reason relations* of implication and incompatibility. The announced goal was to use reason relations in the service of a top-down order of explication that would allow the definition, in terms of them, of the propositional conceptual contents expressed by the declarative sentences that stand in those relations. We have now seen how to fulfill that aspiration, using implications (including those that codify incompatibilities) to define ranges of subjunctive robustness of implications, and then implicational roles, and finally propositional conceptual roles. The top-down explanatory strategy I have pursued was much more thorough-going than that, however. Our aim has been to investigate reason relations of implication and incompatibility in their full generality—not restricting our attention to those that exhibit the topologically and explicitationally closed structure of purely *logical* consequence and inconsistency. That substantial generalization raises a question, though. How do we know that these structurally open, nonmonotonic and nontransitive relations still deserve to be thought of as *reason* relations?

The beginning of a response is to be found in the role of reason relations as specified in the bilateral normative pragmatic metavocabulary whose outlines I sketched. The rest of the answer to that question is to be found in structures at a still higher level: the whole constellation of metavocabularies in terms of which reason relations can be specified as such. In my regimented usage, a 'vocabulary' is a lexicon of sentences together with a set of reason relations defined on them. Those reason relations can be specified in many ways, and any way of doing that is a *rational metavocabulary*.[53] We considered particular, especially perspicuous instances of four kinds of rational metavocabularies in this sense. The bilateral normative pragmatic metavocabulary I sketched characterizes the *use* of declarative sentences in the most minimal recognizable discursive practice. The truth-maker alethic modal semantic metavocabulary characterizes the *meaning* of declarative sentences in terms of the mereologically structured

[53] They are 'metavocabularies' in that they articulate the reason relations of other vocabularies. I think of these rational metavocabularies as themselves vocabularies in the technical sense, but I have not done the work here to show that they can be exhibited in that specific form.

metaphysical states that make them true or false. The sequent-calculus metalogical metavocabulary shows how the reason relations of any base vocabulary can be conservatively extended to include logically complex sentences formed from the sentences of the lexicon of that base vocabulary, and how those logical sentences make it possible to explicitly express reason relations. And the implication-space semantic metavocabulary, as the language of roles, shows how to construct, control, and manipulate propositional conceptual contents on the basis of reason relations among the sentences that express those contents.

The argument is that reason relations are whatever can be specified, elaborated, and discussed by metavocabularies of all these four kinds. Reason relations are the common topic quadrangulated by the metaconceptual perspectives provided by all these kinds of rational metavocabularies. We carefully crafted our versions of each general kind of metavocabulary to ensure that its expressive reach extended to radically substructural or open-structured reason relations in the base vocabularies to which it is applied. So the top-down methodology that seeks to understand propositional contents in terms of reason relations extends to a further upper layer, aiming to understand reason relations in terms of the rational metavocabularies we use to talk about them.

Compelling as (I hope) this response is, we can ask further: in exactly what sense do these very different kinds of metavocabulary afford perspectives on a common topic? Responsibly adapting the visual, spatial metaphor of perspectives to the case of different *conceptual* specifications of one object requires at least a characterization of the space the different perspectives occupy, and a correlation between their 'position' in that space and how things 'look' from that position. Our four metavocabularies are of two different kinds.

Both the bilateral pragmatic metavocabulary and the truth-maker semantic metavocabulary offer substantive *explanations* of reason relations. In the bilateral normative sense, a set of sentences Γ implies a set of sentences Δ just in case commitment to accept all of Γ precludes entitlement to reject all of Δ. In the truth-maker alethic sense, Γ implies Δ just in case every mereological fusion of any truthmaker of all of Γ with any falsity-maker of all of Δ is an impossible state. These are very different explanations. Each appeals to concepts native to its

own setting, which go far beyond what is made available in the spare structure of a vocabulary—namely concepts such as acceptance/rejection, commitment/entitlement, metaphysically possible or impossible states, and their mereological sums. We can say that the bilateral pragmatic and truth-maker semantic rational metavocabularies are *extrinsic-explanatory* rational metavocabularies. They appeal to features *extrinsic* to the reason relations of the vocabularies they address, and do so in order to offer substantive *explanations* of implication and incompatibility.

By contrast, the logical and implication-space metavocabularies are *intrinsic-explicative* rational metavocabularies. They are *intrinsic* in that they appeal to no conceptual resources beyond the minimal formal representation of reason relations as sets of pairs of sentences drawn from a lexicon, that is, resources offered by the abstract concept of a vocabulary. And their aim is not to *explain* what it is for reason relations to obtain between sets of sentences, but to provide the metaconceptual resources to *say* what those reason relations are and *make explicit* the conceptual roles sentences play in those reason relations.

Just as the top-down methodology requires understanding the propositional contents expressed by sentences in terms of their role with respect to reason relations among those sentences, and reason relations in terms of the rational metavocabularies that explain or articulate them, so too that methodology invites us to understand those metavocabularies in terms of their relations to one another. I have emphasized two of these, and mentioned a third in passing. First is the isomorphism at the level of reason relations that Ulf Hlobil crafted between the bilateral normative pragmatic extrinsic-explanatory metavocabulary and the truth-maker alethic modal semantic extrinsic-explanatory metavocabulary. The intrinsic-explicative metavocabularies should both be thought of as making explicit reason relations in the abstract sense that that isomorphism shows to be common to the otherwise disparate extrinsic-explanatory metaconceptual frameworks. Second is the strict correspondence between structural features of the sequent-calculus metavocabulary for computing the reason relations of logically complex sentences from the reason relations of base vocabularies, on the one hand, and the construction of conceptual roles from conceptual roles by the operations of adjunction and symjunction, on the other hand. These two tight structural covariances at the level of reason relations, which

characterize the two orthogonal axes of extrinsic-explanatory and intrinsic-explicative rational metavocabularies, are then cross-connected by a third. For as I reported, for each truth-maker model we can construct an implication-space model that articulates the same reason relations, and for each implication-space model we can construct a truth-maker model that endorses the same reason relations. These relations among the four cardinal varieties of rational metavocabulary functionally characterize the space of positions from which each metavocabulary affords a perspective on reason relations, and (so) propositional conceptual contents.

Here is a final perspective on this intricately structured constellation of rational metavocabularies. Discursive awareness, consciousness in the sense of sapience rather than the mere sentient consciousness manifested in being awake rather than asleep, consists in being able to respond to things by making propositionally contentful claims about them: being able to commit oneself to accept or reject claimable contents and keep track of the way entitlements to such commitments interact and depend on claims offered as reasons for or against others. The bilateral normative pragmatic metavocabulary I sketched in my first lecture accordingly articulates a kind of *theoretical pragmatic self*-consciousness. For it provides the expressive power that enables us theorists to *say* what it is that practitioners must *do* in order thereby to count as making conceptually contentful propositional assertions and denials, and so to be conscious in the sense of sapient: able to claim that things are thus-and-so. The truth-maker rational metavocabulary enables a corresponding sort of *representational semantic self*-consciousness. For it provides the expressive resources to make explicit crucial relations between claimable conceptual contents and the worldly states that make them true or false.

Base vocabularies that have been extended and expressively enriched by the introduction of *logical* vocabulary are the organs of a distinctive kind of *pure* rational self-consciousness. For, as we have seen, sentential logical vocabulary embodies the expressive power to make explicit reason relations of implication and incompatibility, both nonlogical and logical. Intrinsic implication-space rational metavocabularies are the organs of a distinctive kind of pure *semantic* theoretical self-consciousness. For they are the native language for specifying the conceptual roles sentences play in virtue of standing to each other in reason relations. So the isomorphisms

and correspondences among the four different kinds of rational metavocabulary I have been rehearsing as the very top-most level of the top-down order of explication sketched here can be understood as articulating the internal structure of rational self-consciousness as such. I am happy to close these lectures bathing in the Hegelian resonances of that thought.

End of Lecture III