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 Boolean aggregative 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.'

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 elated 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.[1] 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

[1] 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.[2] 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.[3]

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

[2] 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].

[3] "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.[4] 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

[4] 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: $\forall X,Y \subseteq L[\ RSR_M(\langle X,Y\rangle) =_{df.} \{\langle W,Z\rangle : W,Z \subseteq L \text{ and } \langle X \cup W, Y \cup Z\rangle \in \mathbf{I}_M\}\ ]$.[5]

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).[6] We accordingly move from understanding the subjects of semantic interpretation to be

[5] 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.

[6] 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*).[7]

The premissory and conclusory roles are *different* sets of implications.[8] 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.

[7] 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.
[8] 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*, $<\Gamma,\Delta>$[9] 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

[9] 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.[10] 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).[11]

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

[10] 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.

[11] 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.[12] 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

[12] 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 $\langle\Gamma,\Delta\rangle$.[13]

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.[14]

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

---

[13] 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.

[14] 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'.

To address that question, we need 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 idea is that where the implicational role of A, [A] is $\langle a^+,a^-\rangle$, we can compute the implicational role of a compound f(A) by $[f(A)] = \langle a^-,a^+\rangle$. It is not hard to recognize the operator *f* that is so defined as *negation*. 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 ¬A is added as a premise—that, is it will be in the premissory implicational range of ¬A. The metainferential rules that relate the occurrence of A on one side of the turnstile to ¬A on the other are equivalent to defining the implicational role of ¬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. 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:

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

Each of those elements can be further decomposed:

$\mathbf{[A] = \langle a^{+}, a^{-} \rangle = \langle \mathcal{R}(\{\langle A, \varnothing \rangle\}), \mathcal{R}(\{\langle \varnothing, A \rangle\}) \rangle}$.

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:

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

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

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

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

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→B, which is

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

And the corresponding left rule of the sequent calculus.

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

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→. 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.[15] 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).[16]

[15] *RLLR* 5.2.2, pp. 224 ff.
[16] *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.[17] 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.[18]

---

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

[18] 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.[19]

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.

[19] 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.[20]

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

[20] 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*.[21] 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

[21] 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