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 <u>reason relations</u> 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*.[1] 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;

[1] 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 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.[2] 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

---

[2] 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.[3] 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.[4]

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 designated 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 metametaconceptual relational features of bivalent truth values that explain their metaconceptual capacity to express

[3] 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.

[4] 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*.

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 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 the two 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.[5]

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?

---

[5] 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.[6] As with Dummett's notion of harmony, 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

[6] 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.[7] 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.

[7] 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.[8]

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:

[8] 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): $\dfrac{\Gamma |\!\sim A \quad \Gamma |\!\sim B}{\Gamma, B |\!\sim A.}$

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.[9]

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.[10] 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

[9] 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.

[10] 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):

$$\frac{\Gamma, B \mid\!\sim A \qquad \Gamma \mid\!\sim B}{\Gamma \mid\!\sim A.}$$

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 hypernonmonotonic, 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 <u>reason relations</u> is called upon to carry out. For the logical 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).[11]

[11] 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.[12] 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

[12] 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): Γ, A |~ B, Δ

========= Bidirectional Inference Line

Γ |~ A→B, Δ.

(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): Γ, A |~ Δ

========= Bidirectional Inference Line

Γ |~ ¬A, Δ.

(This is the multisuccedent version of Γ |~ ¬A if and only if Γ#A, i.e. Γ,A |~ .)

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.[13]

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.

[13] 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.[14]

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.[15]

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

---

[14] 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*.

[15] 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.

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.[16] 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 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.

[16] (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.

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¬: $$\frac{\Gamma \mid\!\sim \Delta, A}{\Gamma, \neg A \mid\!\sim \Delta}$$

R¬: $$\frac{\Gamma, A \mid\!\sim \Delta}{\Gamma \mid\!\sim \Delta, \neg A}$$

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

R→: $$\frac{\Gamma, A \mid\!\sim B, \Delta}{\Gamma \mid\!\sim A \rightarrow B, \Delta}$$

L&: $$\frac{\Gamma, A, B \mid\!\sim \Delta}{\Gamma, A \& B \mid\!\sim \Delta}$$

R&: $$\frac{\Gamma \mid\!\sim \Delta, A \quad \Gamma \mid\!\sim \Delta, B \quad \Gamma \mid\!\sim \Delta, A, B}{\Gamma \mid\!\sim \Delta, A \& B}$$

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

R∨: $$\frac{\Gamma \mid\!\sim \Delta, A, B}{\Gamma \mid\!\sim \Delta, A \vee B}$$