October 23, 2024

Notes for Meeting 7

Implication-Space Semantics: The Pure Theory of Conceptual Roles

Recap:

- I surveyed two ways of thinking about the relations between *reasons* and *logic*: *logicism*, which takes all genuinely *good* reasons to be at base *logically* good reasons, and *expressivism*, which understands the job of logic not to be *determining* what are good reasons but giving us the expressive power to *say* what are good reasons.
- The expressivist project is to show how to extend arbitrary base vocabularies, adding new sentences formed from the old by the application of logical operators, and computing the reason relations governing the extended language from those of the base vocabulary. In addition to being conservatively *elaborated from* the reason relations of the base vocabulary, suitable logical vocabulary must make it possible to *specify*, that is, *say explicitly* (in the form of claimables), in the logically extended vocabulary, what reason relations hold in that vocabulary, including the base. These criteria of adequacy are summed in the phrase that we want a logic that is *universally LX*: *el*aborated from and *ex*plicative of any base vocabularies whatsoever.
- Sequent-calculus metavocabularies (following Gentzen) were introduced as ways of formulating metainferential rules that derive reason relations for more logically complex sentences from reason relations for logically simpler sentences.
- The logic NMMS (for Non-Monotonic MultiSuccedent), was introduced, and shown to be universally LX, in part by a powerful expressive representation theorem due to Dan Kaplan:

  > For any set *AtomicImp* of sequents in any base vocabulary L_B, there is a nonempty set of sequents *ExtImp* in the logical extension of L_B by NMMS such that every individual sequent in *ExtImp* is derivable if and only if all the sequents of *AtomicImp* hold. And conversely, for any set of sequents *ExtImp* in the logical extension of L_B by NMMS, we can compute the exact set *AtomicImp* of sequents defined on the lexicon of the base vocabulary that must hold in order for all the sequents in *ExtImp* to be derivable.

This remarkable result means that if we pick any arbitrary sequent (or set of them) in any vocabulary logically extended by applying the rules of NMMS, **we can compute exactly what that sequent (or set of them) *says* about the base from which it is elaborated: it will be a *good* sequent just in case this particular set of atomic implications holds in the base**. This is what all bases have in common that validate the initial target logically complex sequent or sequents.

Now one might say: “That isn’t so remarkable. I can do the same thing in a two-valued classical setting. The truth tables let me work back from complex sequents or the corresponding candidate theorems with conditionals (since classical logic is compact).”
True. And remember, in a clear sense, NMMS just *is* classical logic: that’s the logic its rules determine in the context of CO, MO, and Mixed-Context Cut.
What we are seeing is a generalization of this nice feature of classical two-valued logic.
But Dan’s expressive completeness result holds even for radically substructural bases, which in order to extend open-structured bases *conservatively* must define open-structured reason relations among sets of logically complex sentences, too. And no analogue of truth-tables degrades gracefully in such an unstructured setting.

- At the end I showed how the familiar three-valued logics K3 and LP show up in an inferentialist-expressivist setting as the logics of premissory and conclusory role inclusions, respectively.

The Hlobil pragmatic-semantic isomorphism.
We can think about what is in common, between items mapped onto one another: roles w/res to the reason relations that the mapping preserves.

I have suggested that we can think of the pragmatic and semantic MVs as vocabularies: sets of reason relations on a lexicon.

Now we get to see what we get from following out that suggested simplified specification (that vocabulary MV) or model—and so the source of my retrospective confidence in this vocabulary model. What we get is an abstract, tractable specification of *conceptual roles*.

We're trying to specify reason relations articulating a vocab for specifying reason relations. And there's a certain equilibrium or harmony required in doing that.

I. Implication-Space Semantics

Implication-space models just are vocabularies.
Implication spaces, the basic model-theoretic structures of the Kaplan-Hlobil semantics, are just what I have been calling "vocabularies." For vocabularies consist of a *lexicon* and a set of *reason relations* among sets of lexical items. And implication space models consist of a set of candidate implications, which is determined by and computable from a lexicon, and a distinguished subset of that larger space, thought of as the *good* implications, which (since they also encode incompatibilities) just is the *reason relations* the implication-space model validates. Implication-space model-theoretic semantics just fills in and makes more determinate both elements of vocabularies.

a) The Hlobil isomorphism between reason relations construed pragmatically (in bilateral, normative terms) and construed semantically (in truth-maker model-theoretic terms) shows that if collections of assertions and denials of atomic sentences are normatively ruled out just in case all fusions of truth-makers of the asserted sentences and falsity-makers of the denied sentences are alethically ruled out, then the pragmatic-normative theory of consequence among sentences specified in a normative pragmatic MV coincides with the semantic-representationalist truthmaker theory of consequence among worldly propositions.
This fact raises a question and a challenge:
Q: Can we independently specify those shared roles with respect to reason relations specified in the two kinds of metavocabulary? That is, is there a rational metavocabulary that will specify the common, shared, roles sentences play with respect to reason relations specified in the pragmatic and representational semantic MVs, abstractly, without commitment to anything that differentiates the pragmatic and semantic MVs?
A: Yes. The implication-space semantics that Dan Kaplan adapted and developed based on Girard's phase-space semantics for linear logic, and which Ulf Hlobil then adapted and developed based on Dan's presentation in his dissertation. It gives us a way to understand

the *rational forms*, the *conceptual roles*, sentences play in virtue of standing to one another in reason relations among sets of sentences, which we have seen can be specified more concretely in the two contrasting pragmatic and semantic rational MVs.

Implication-space semantics is an attempt to do *model-theoretically* what the sequent calculus does *proof-theoretically*. Namely, it seeks to *take reason relations* (represented by set-theoretic constructions on sentences or other bearers) *as objects*, on which to perform operations and among which to explore relations. Any semantic model theory must determine reason relations (else it is not doing semantics), but usually it does so by looking at relations among some *other* structures.

The substantial, notorious expressive advantages of model-theoretic over proof-theoretic rational metavocabularies is that model-theoretic MVs typically have the full expressive power of iterated alternating quantifiers, formulating semantic conditions of forms such as "For any object x, if Px then there exists an object y s.t. Rxy. The MV of the sequent calculus does not permit this. Indeed, we have to stretch and extend the sequent-calculus MV so as to distinguish *persistent* sequents: ones that are good, and remain good on arbitrary addition of premises or conclusions. And that is just one quantifier.

<u>Three Results:</u>

1. We will see that **implication-space semantics (ISS) permits a sound and complete semantics for the universally LX logic NMMS**.
2. Also: can translate truth-maker models into implication-space models, and *vice versa*.
3. But we will also see something more remarkable, something that bears comparison with the Hlobil isomorphism at the level of reason relations between (the right sort of) bilateral normative pragmatic MV and (the right sort of) truth-maker semantics.

That is an **isomorphism between (the right sort of) proof-theoretic (sequent-calculus) MV** for specifying and manipulating reason relations **and (the right sort of) model-theoretic (implication space) MV** for specifying and manipulating reason relations.
This isomorphism is far more general, extending well beyond NMMS to offer model-theoretic semantic codification of any sequent rules whatsoever (e.g. those for **linear logic**), and conversely, proof-theoretic codifications of a broad range of model-theoretic conditions.
Our result relating sequent-calculus proof-theoretic MVs and implication-space model-theoretic MVs goes well beyond any previous gestures in this direction, offering a robust and general correlation that can be exploited in either direction.

A *vocabulary* V = <L,R> is i) a lexicon L together with ii) a specification R of *reason relations* defined on that lexicon. (i) is just a set of sentences. (ii) can be a set of ordered pairs <X,Y> of sets of sentences of L, where $<X,Y>\in R$, for X,Y subsets of L and Y is nonempty means that X|~Y is a good (multisuccedent) implication, and if $Y=\emptyset$, then $<X,Y>\in R$ means that X is incoherent (so any subset of it is incompatible with the remainder).

An *implication* ***space*** *S* is the set of all *candidate implications* defined on a set of bearers, for instance, a lexicon L, as the set of all ordered pairs of sets of bearers: $S = \mathcal{P}(L) x \mathcal{P}(L)$.
An *implication* ***frame*** <*S*, **I**>is an implication space together a set of distinguished candidate implications $\mathbf{I} \subseteq S$, interpreted as the *good* implications.
Implication frames are really just encodings of vocabularies, in our technical sense.

In his presentation of implication spaces in Chapter Five, Ulf is scrupulously careful to refer to the elements of the underlying language or lexicon as "bearers" rather than as "sentences." That is because when the bearers of sentential conceptual roles are specified in the semantic truthmaker MV, they are not *sentences*, but *worldly propositions*, pairs of sets of truthmaking states and sets of falsity-making states—*some* of which can then be associated with the sentences of some language-or-lexicon by a semantic *interpretation* function. His point is important, and his care is appreciated, but (partly looking forward to the characterization of implication-space vocabularies as *intrinsic* rational metavocabularies) I will talk about "sentences" rather than the more general and more correct "bearers."
The "space" in "implication space" is the set of all *candidate implications* that can be formed from the underlying language or lexicon.
We think of these as pairs of sets of sentences of the lexicon.
The first element is a set containing the premises and the second element is a set containing the conclusions of the candidate implication.
We say "*candidate*" implication because there is no question yet about whether the elements of our implication space are *good* implications. That is settled by the second element of implication spaces.
So the space S of an implication space defined on the lexicon L is $\mathcal{P}(L) x\ \mathcal{P}(L)$, the cross-product of the powersets of the lexicon.

Step One: Semantically interpret candidate *implications*

The idea that there is such a thing, and that it should be addressed *first*, before we worry about the semantic interpretation of *sentences*, is one of Dan Kaplan's big conceptual innovations (in his dissertation).

Speaking of Kant turning on its head the traditional bottom-up order of discursive explanation, which proceeds from a **doctrine of *concepts*** (particular and general), erects on it a **doctrine of *judgments*** (judgeables), and then finishes with a **doctrine of *consequences***, in the form of syllogisms, Sellars says:

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

Elsewhere (CDCM) Sellars argues that to count as *descriptions*, and not just *labels*, that is, to have specifically *conceptual* content, linguistic expressions must not just have appropriate *circumstances* of application, but also must be "**situated in a space of implications**", which means that they have inferential appropriate *consequences* of application.

Sufficiently knowing one's way practically about in that "space of implications" requires not only having a good enough rough-and-ready practical capacity to distinguish materially good from materially bad implications (and incompatibilities), however partial and fallible, but also having a good enough rough-and-ready practical capacity to distinguish, for each candidate material implication, a ***range of subjunctive robustness***: that is, a practical capacity to distinguish additional circumstances that (as additional premises) would or would not *infirm* the implication, that would *keep* it or *make* it a materially good one.
These correspond to the ***extension* and *intension*** of *goodness*-values of candidate implications.

The bivalued *extension* of a candidate implication <X,Y> is its goodness value,
that is, whether or not <X,Y>∈**I**, meaning that it is a *good* implication, one that holds.
The *intension* of a candidate implication <X,Y> is its *range of subjunctive robustness*.

Being a *good* implication, or not, is the semantic *extension* of an implication, its goodness-value (compare: truth value). The semantic *intension* of an implication is its range of subjunctive robustness: the circumstances under which it *would* be or *would* remain good.
"The idea behind ranges of subjunctive robustness is that, especially in a nonmonotonic setting, it is interesting to consider which additions to an implication do not defeat the implication, if the implication is already good, or turn the implication into a good implication, if it is not already good." [213]
The *range of subjunctive robustness* of a candidate implication <X,Y> is the set of other implications that, when adjoined to <X,Y>, yield a *good* implication, that is, an implication that is in **I**.

RSRs determine *implicational roles*, i.e., the roles of implications w/res to other implications.

The *range of subjunctive robustness* RSR(<X,Y>) of a candidate implication is the set of pairs of sets of sentences that can be added to <X,Y> to *keep* it good (if it is in **I**) or *make* it good (if it is not in **I**):
RSR(<X,Y>) $=_{\text{df.}}$ {<W,Z>∈*S*: <X∪W,Y∪Z>∈**I**}.
Note: The RSR of a *set* of implications is the intersection of the RSRs of the elements of the set.
It follows that **<X,Y>⊆I iff <∅,∅>∈RSR(<X,Y>)**.
So implicational *intension* determines implicational *extension*, without further information.

Candidate (sets of) implications are *implicationally equivalent* iff they have the same RSRs:
G≈F iff RSR(G) = RSR(F), for G,F⊆*S*.
Implicational roles **R(G)** of (sets of) implications are their *implicational equivalence classes*:

R(G)={H⊆*S*: H≈G}.

This is the crucial act of *abstraction*, which lets us express in implication-space terms what is common to the otherwise very different conceptions of reason relations of the pragmatic and semantic MVs related by the Hlobil isomorphism. Abstraction—as in the original Fregean example (from the *Grundlagen*) of moving from lines, via the relation of being parallel, to directions of lines—is defining equivalence classes of items.
Henceforth we don't care whether implications and incompatibilities are understood in deontic normative terms of preclusion from entitlement to some set of joint commitments to accept or reject, or in alethic modal terms of the results of some fusions of states being impossible states. For we are extracting the roles with respect to ranges of subjunctive robustness of implications.

To do any real work, though, we must define operations on or relations among *implicational roles*, so defined. Here, everything turns on picking just the right opearations to look at. Those we use are the results of lots of work over a long time.
(I did none of it. Dan and Ulf figured out what was needed, inspired by Girard.)

We define two operations on implicational equivalence classes:
⊔: **adjunction** R(<X,Y>)⊔R(<W,Z>) $=_{df.}$ R(<X∪W, Y∪Z>).
⊓: **symjunction** R(<X,Y>)⊓R(<W,Z>) $=_{df.}$ R(<X,Y> ∪ <W,Z>).
for <X,Y>,<W,Z>∈*S*, and with corresponding definitions for adjoining and symjoining *sets* of implicational ≈-classes.
As their symbols indicate, *ad*junction is an analogue in this setting of set-theoretic *union*, and *sym*junction is an analogue in this setting of set-theoretic *intersection* (see Note above re RSR).

Symjunction then appeals to RSRs to define an operation on the RSRs of implications.
Definition 69 (Symjunction, ⊓). Let $F, G \subseteq \mathbf{S}$, then:
$R(F) \sqcap R(G) = R(F \cup G)$.
⊓
"The implicational role of a union of sets of candidate implications is, therefore, the sets of candidate implications that share their range of subjunctive robustness with the intersection of the ranges of subjunctive robustness of the sets of which we take the union. Hence, all of the things that yield a good implication when combined with something that plays the role *R*(*F*), and also when combined with something that plays the role *R*(*G*), yield a good implication when combined with something that plays the role of their symjunction *R*(*F*) ⊓ *R*(*G*)." [218]

<u>Step Two: Semantically interpret *sentences* of the lexicon in terms of the semantics of implications in which they appear as premises or conclusions</u>

The semantic interpretation of sentences is downstream conceptually and explanatorily from the semantic interpretation of implications.

i. The idea of the contribution the occurrence of a sentence A makes to the goodness of implications in which it appears as a *premise*, and the potentially different (remember when premissory role inclusions and conclusory role inclusions come apart in nontransitive settings, from the discussion of the trilogics K3 and LP last time) contribution the occurrence of a sentence A makes to the goodness of implications in which it appears as a *conclusion*. Represent the *conceptual role* of a sentence (w/res to the reason relations of a vocabulary in whose lexicon it is) as the ordered pair of its *premissory role* and its *conclusory role*.

ii. The question then is how to make the vague talk of the "*contribution*" the occurrence of the sentence as a premise or conclusion makes to the goodness of implications.
We'll do that by using the notion of the <u>range of subjunctive robustness</u> of implications.

Q: Which implications?

A: <{A},∅> is a candidate implication such that every candidate implication in which A occurs as a premise is the result of adjoining some other candidate implication to it. **The *good* implications in which A occurs as a premise are then just the *range of subjunctive robustness* of <{A},∅>, RSR<{A},∅>**.

Similarly, <∅,{A}> is a candidate implication such that every candidate implication in which A occurs as a conclusion is the result of adjoining some other candidate implication to it. **The *good* implications in which A occurs as a conclusion are then just the *range of subjunctive robustness* of <{A},∅>, RSR<∅,{A}>**.

"Notice that the set of pairs $\langle X,Y\rangle$ such that $\langle X \cup \{\phi\},Y\rangle \in \mathbf{I}$ is the range of subjunctive robustness of $\langle\{\phi\},\emptyset\rangle$. And the set of pairs $\langle X,Y\rangle$ such that $\langle X,Y \cup \{\phi\}\rangle \in \mathbf{I}$ is the range of subjunctive robustness of $\langle\emptyset, \{\phi\}\rangle$. So it follows from what we said above **that these ranges of subjunctive robustness fully determine the premisory and conclusory roles of a given bearer ϕ**."[213]

"To illustrate, the premisory role of bearer ϕ is the set of sets of candidate implications with the same range of subjunctive robustness as $\langle\{\phi\},\emptyset\rangle$. For example, the singleton of a candidate implication $\{\langle\Gamma, \Delta\rangle\}$ is in this set if and only if parallel additions of premises or conclusions to $\langle\Gamma, \Delta\rangle$ and to $\langle\{\phi\},\emptyset\rangle$ always either both yield good implications or both yield implications that are not good. In other words, in any candidate implication, we can replace the bearer ϕ as a premise *salva consequentia*— that is, without turning a good implication into a bad one or vice versa— with the combination of Γ as premises and Δ as conclusions, and the other way around. **So, if two bearers have the same premisory role, then they can be substituted for each other**

as premises *salva consequentia*. Conclusory roles are analogous, except that $\langle\{\phi\},\emptyset\rangle$ is changed to $\langle\emptyset, \{\phi\}\rangle$. Hence, if two bearers have the same conclusory role, they can be substituted for each other as conclusions *salva consequentia*. The equivalence classes with respect to ranges of subjunctive robustness that are implicational roles thus capture the idea of "playing the same role in implications."

According to the notion of an implicational role that we have just introduced, what it means to play a particular implicational role is to be a member of a particular equivalence class of things with the same range of subjunctive robustness. **For the case of individual bearers of implicational roles, their roles are pairs of such equivalence classes, namely the premissory and the conclusory roles of the bearer**." [215]

If A is a sentence in the lexicon L, then $\mathbf{R^+(A) =_{df.} R(\langle\{A\},\emptyset\rangle)}$ is the implication equivalence class of all good implications in which A appears as a *premise*. And $\mathbf{R^-(A) =_{df.} R(\langle\emptyset,\{A\}\rangle)}$ is the implication equivalence class of all good implications in which A appears as a *conclusion*. The pair of these is the implicational role of sentence A, $\mathbf{R(A) =_{df.} \langle R^+(A), R^-(A)\rangle}$.

We can also form a class **C** of possible *conceptual contents*, which might or might not be the implicational roles of any sentences of L. If F, G $\subseteq$ S are *any* sets of candidate implications, $\langle R(F),R(G)\rangle$ is a conceptual content, which we can call $a \in \mathbf{C}$, so that $\mathbf{R(a)=\langle a^+ ,a^- \rangle = \langle R(F),R(G)\rangle}$.

Q: What kinds of monsters are there among the "conceptual contents" so defined?
We see below that "tonk" sentences are.
Have we at all explored the various sorts of pathologies that there can be among conceptual contents? (I'm thinking of 'tonk'-like combinations $\langle a^- ,a^- \rangle$ and $\langle a^+ ,a^+ \rangle$.
There might well be good stuff to be found here.
Both $\langle\emptyset,\emptyset\rangle = e$ is here, and $\langle L,L\rangle$.
There might be plausible structural principles that could be imposed here.

<u>Interpretation Function:</u>
[.] maps sentences of the language L to some conceptual contents **C**.
The logical connective clauses that an interpretation must respect are:

If $A \in L$ is an atomic sentence, then $[A] =_{df.} \langle a^+, a^-\rangle \in \mathbf{C}$.

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

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

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

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

Exercise: Define 'tonk' semantically in these terms.

Response: It should have the *intro* (left) rule of *dis*junction and the *elimination* (right) rule of *con*junction. These correspond to the first and last elements of the roles, so:
[AtonkB] = $\langle a^+ \sqcap b^+ \sqcap (a^+ \sqcup b^+), a^- \sqcap b^- \sqcap (a^- \sqcup b^-)\rangle$.

Explanation:

Negation:

"Negation swaps the premisory and conclusory roles, so that the content that an interpretation function assigns to a negation $\neg\phi$ is always the content that is like the content that the function assigns to ϕ except that the premisory and the conclusory roles are swapped.
Hence, the ranges of subjunctive robustness of $\langle\{\neg\phi\},\emptyset\rangle$ and $\langle\emptyset, \{\neg\phi\}\rangle$ are, respectively, the ranges of subjunctive robustness of $\langle\emptyset, \{\phi\}\rangle$ and $\langle\{\phi\},\emptyset\rangle$. So, whatever yields a good implication when combined with $\langle\{\neg\phi\},\emptyset\rangle$ also yields a good implication when combined with $\langle\emptyset, \{\phi\}\rangle$.
This is how negation allows us to make explicit incompatibility. For we encode the fact that a set of bearers Γ is incompatible with the bearer ϕ in our implications as $\langle\Gamma \cup \{\phi\},\emptyset\rangle \in \mathbf{I}$. And our negation clauses ensure that we can substitute the negation for the negatum on the other side *salva consequentia* to get: $\langle\Gamma, \{\neg\phi\}\rangle \in \mathbf{I}$.
Thus, **the negation makes explicit, on the right side of an implication, that it is ruled out by what is on the left side. This is precisely the explicative potential that we described as the essential function of negation: the II Incoherence-Incompatibility property**." [218-19].

Conditional:

"**The content of a conditional always has as its conclusory role the adjunction of the premisory role of the antecedent and the conclusory role of the consequent.**
That means that the range of subjunctive robustness of $\langle\Gamma, \Delta \cup \{\phi \rightarrow \psi\}\rangle$ is always identical to the range of subjunctive robustness of $\langle\Gamma \cup \{\phi\}, \Delta \cup \{\psi\}\rangle$.
This ensures that the conditional makes implications explicit in the way described in Chapter Three. Namely, it ensures that $\langle\Gamma \cup \{\phi\}, \Delta \cup \{\psi\}\rangle \in \mathbf{I}$ just in case $\langle\Gamma, \Delta \cup \{\phi \rightarrow \psi\}\rangle \in \mathbf{I}$, which is a formulation of the Deduction-Detachment Condition on Conditionals from Chapter Three.

And **the premisory role that an interpretation function assigns to a conditional is the symjunction of the conclusory role of the antecedent, the premisory role of the consequent, and the adjunction of these roles.**
That means that the range of subjunctive robustness of $\langle\Gamma \cup \{\phi \rightarrow \psi\}, \Delta\rangle$ is always the intersection of the range of subjunctive robustness of $\langle\Gamma, \Delta \cup \{\phi\}\rangle$ and the range of subjunctive robustness of $\langle\Gamma \cup \{\psi\}, \Delta\rangle$ and the range of subjunctive robustness of $\langle\Gamma \cup \{\psi\}, \Delta \cup \{\phi\}\rangle$.

In this way, the semantic clauses for the conditional assign a conceptual content to conditionals by placing the use of the conditional as a premise and as a conclusion in equivalence classes with respect to ranges of subjunctive robustness, that is, by assigning the conditional its premisory and conclusory implicational role. The clauses for conjunction and disjunction can be understood in an analogous way." [219]

Completeness Theorem: This semantics in terms of implication spaces and frames is *sound and complete* for the logical vocabulary of NMMS, no matter what base vocabulary both are elaborated from.
If in the base vocabulary, every implication $\langle X,Y\rangle$: $X\cap Y\neq\varnothing \Rightarrow \langle X,Y\rangle\in\mathbf{I}$, then CO holds in the base and also in its logical extension by NMMS, and the logical extension is *supraclassical*.
If the converse also holds, $\langle X,Y\rangle\in\mathbf{I} \Rightarrow X\cap Y\neq\varnothing$ in the base vocabulary, the pure logic that holds in all those models is just *classical logic*.

It is easy to extend these techniques to give rules for *linear logic* (Multiplicative and Additive Linear Logic = MALL), as Ulf does in Section 5.4.2. This is not at all surprising given that Dan Kaplan adapted and built on Girard's phase-space semantics for linear logic to come up with ISS.

<u>Correspondence of sequent calculus *proof* theory and implication space *model* theory</u>:

"The semantic clauses for the logical constants in implication-space semantics correspond to the sequent rules of NMMS in a perhaps surprisingly direct way. To see this consider the rules for the conditional, as an example:

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

R→: $$\frac{\Gamma, A|\sim B, \Delta}{\Gamma \mid\sim A\to B, \Delta}$$

Since these rules are invertible, the right rule implies that $\Gamma|\sim A\to B, \Delta$ is derivable if and only if $\Gamma, A|\sim B, \Delta$ is derivable. This is the Deduction-Detachment Condition from last time.
It corresponds in implication-space semantics to the claim that $\Gamma|\sim^{M} A\to B, \Delta$ holds in a model M if and only if $\Gamma,A|\sim^{M} B, \Delta$ holds in that model.
Now, $\Gamma|\sim^{M} A\to B, \Delta$ holds just in case any implication that plays the role of $\langle\Gamma, \Delta\rangle$ yields a good implication when it is combined with anything that plays the conclusory role of $A \to B$.
So $\{\langle\Gamma, \Delta\rangle\}$ must be in the range of subjunctive robustness of $\langle\varnothing,\{A\to B\}\rangle$, and the same must hold for any other implications with the same role as $\langle\Gamma, \Delta\rangle$.
Our semantic clauses for the conditional tell us that the role of $\langle\varnothing,\{A\to B\}\rangle$ is $a^{+}\sqcup b^{-}$, that is, it is the role of $\langle\{A\},\{B\}\rangle$.
But two things with the same role share their ranges of subjunctive robustness.
So everything with the role $\{\langle\Gamma, \Delta\rangle\}$ is in the range of subjunctive robustness of $\langle\varnothing,\{A\to B\}\rangle$ just in case it is in the range of subjunctive robustness of $\langle\{A\},\{B\}\rangle$.
So, $\Gamma,A|\sim^{M} B, \Delta$.

For the left rule, [L→], the analogous relation holds. Because the rule is invertible, it tells us that Γ,A→B|~Δ is derivable if and only if all of Γ|~A, Δ and Γ,B|~Δ and Γ,B|~A, Δ are derivable.
Suppose that $\Gamma, A \to B |\sim^M \Delta$.
Then the premisory role of *A* → *B* is such that combining anything that plays that role and anything with the role of <Γ, Δ> yields a good implication.
Hence, <Γ, Δ> is in the range of subjunctive robustness of <{A→B}, ∅>.
The semantic clause for the conditional informs us that the role of <{A→B}, ∅> is
a−⊓b+⊓ (a−⊔b+).
And the range of subjunctive robustness of a symjunction of roles is the intersection of the ranges of subjunctive robustness of the roles whose symjunction it is.
So, the range of subjunctive robustness of $a^- \sqcap b^+ \sqcap (a^- \sqcup b^+)$ is the intersection of the ranges of subjunctive robustness of <∅,{A}> and <{B},∅> and {{B},{A}>.
Therefore, <Γ, Δ> is in the ranges of subjunctive robustness of all of <∅,{A}> and <{B},∅> and {{B},{A}>.
And that is just another way to say that $\Gamma |\sim^M A, \Delta$ and $\Gamma, B |\sim^M \Delta$ and $\Gamma, B |\sim^M A, \Delta$ all hold." [223]

"So, the rules of NMMS are not only equivalent to the semantic clauses of truth-maker theory…, but they are also equivalent to the semantic clauses of implication space semantics.
Indeed, we can formulate this correspondence in a general way as follows.

- 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**." [*RLLR* 223]

<u>NMMS</u>:

L→: Γ|~Δ, A   Γ, B|~Δ   Γ, B|~Δ, A
====================
Γ, A→B |~ Δ

R→: Γ, A|~ Δ, B
=========
Γ |~ Δ, A→B

L¬:	$\begin{array}{c} \Gamma \mid\!\sim \Delta, A \\ \hline\hline \Gamma, \neg A \mid\!\sim \Delta \end{array}$	R¬:	$\begin{array}{c} \Gamma, A \mid\!\sim \Delta \\ \hline\hline \Gamma \mid\!\sim \Delta, \neg A \end{array}$
L∧:	$\begin{array}{c} \Gamma, A, B \mid\!\sim \Delta \\ \hline\hline \Gamma, A \wedge B \mid\!\sim \Delta \end{array}$	R∧:	$\begin{array}{c} \Gamma \mid\!\sim \Delta, A \quad \Gamma \mid\!\sim \Delta, B \quad \Gamma \mid\!\sim \Delta, A, B \\ \hline\hline \Gamma \mid\!\sim \Delta, A \wedge B \end{array}$
L∨:	$\begin{array}{c} \Gamma, A \mid\!\sim \Delta \quad \Gamma, B \mid\!\sim \Delta \quad \Gamma, A, B \mid\!\sim \Delta \\ \hline\hline \Gamma, A \vee B \mid\!\sim \Delta \end{array}$	R∨:	$\begin{array}{c} \Gamma, \mid\!\sim \Delta, A, B \\ \hline\hline \Gamma \mid\!\sim \Delta, A \vee B \end{array}$

This isomorphism of the model theory with proof theory means we have general rules for constructing the implication-space semantics for any logical connective definitions we can specify in the sequent-calculus metavocabulary.

Intuitionism uses single-succedent sequents. There is an analogous implication-space construction, whose candidate implications are pairs of a set of premises and a singleton conclusion. Reversible rules are not as generally available in the single-succedent setting, so no close analogue of Kaplan's expressive representation theorem is provable here.
We have not explored single-conclusion implication spaces much at all.
There are probably good things to be discovered.

Relation of Implication-Space Semantics to Truthmaker Semantics on State Spaces:

Define implication-space frames from modalized truthmaker state spaces:
"We can define an implication frame for any modalized state space, <S, $S^{\lozenge}$, ⊑>, by letting the bearers be worldly propositions, that is, pairs of sets of states from *S*, and defining the good implications by appeal to impossible states." [224]

"It follows from Theorem 79 that if there is a truth-maker model in which exactly a particular set of implications holds among the interpreted sentences, then there is an implication-space model in which exactly the same implications hold. The theorem ensures that for every truthmaker model, there is a parallel implication frame model such that the consequence relation defined by these models coincide." [226]

Define modalized truthmaker state spaces from implication-space models:

"We can also go in the other direction. If we are given an implication-space model that codifies a particular consequence relation over a language, then we can construct a truth-maker model that codifies the same consequence relation over that language." [227]

Let the set S of states be the set of pairs of sets of sentences of the (atomic) lexicon, <X,Y> for X,Y⊆L. Define the mereological part-whole relation among states (from which fusion of states is defined) ⊑ by <X,Y>⊑<W,Z> iff (X⊆W and Y⊆Z). <X,Y> is a *possible* state just in case <X,Y>∉**I**. Worldly propositions are pairs of sets of states (potential truth-makers and falsity-makers)—which are interpreted here as pairs of sets of pairs of sets of atomic sentences—satisfying Fine's Exclusivity. Consequence relations among pairs of sets of worldly propositions are defined as Hlobil does: as holding when every result of fusing any truthmaker of all of the premises with any falsity-maker of all of the conclusions is an *im*possible state. It is shown that these are exactly those determined by **I** in the original implication-space model.

[**Read these passages in class**—or at least discuss these points:]
"It follows immediately from this theorem that truth-maker theory and implication-space semantics are equivalent in their power to provide counterexamples to implications. Hence, if we think of model theory as a way to provide counterexamples to implications, then these two theories are equivalent as model theories. However, we have seen that implication space semantics provides an account of conceptual contents as roles in implications, while truth-maker theory allows for different worldly propositions that play the same implicational roles. In this sense, the interpretants of implication-space semantics are more abstract; they correspond to equivalence classes of the interpretants of truth-maker theory." [227]

"Let us take stock. We can use implication-space semantics to interpret the sequent calculus NMMS and truth-maker theory in parallel ways. Both theories provide us with a space of bearers of implicational roles. These are sentences in NMMS and they are worldly propositions in truthmaker theory.

Moreover, both theories provide us with an implication relation over sets of bearers; this relation is |~ in NMMS and $\mathbf{I}_{\diamond}$ (defined by $S^{\Diamond}$) in truth-maker theory. So, we can abstract conceptual contents from both theories. We can give accounts of the logical connectives in all three theories, and they have exactly corresponding effects in their respective theories. That is, the operational rules of NMMS and the semantic clauses for the connectives in truth-maker theory are both equivalent to the semantic clauses of implication-space semantics. So, the three theories are equivalent as logical theories.

However, implication-space semantics is, in one sense, the most abstract theory of the three theories. For it is an *intrinsic* semantic theory: it uses reason relations itself to interpret the items that stand in these reason relations. The codification of reason relations in the metavocabulary of implication-space semantics does not explain these reason relations in terms of normative exclusion relations among discursive acts or alethic exclusion relations among worldly states. Implication-space semantics does not specify what the bearers of implicational roles are or what the nature of exclusion relations between them is. We have thus arrived at an intrinsic characterization of the rational forms that we encountered in NMMS and in truth-maker theory." [228]

Premissory and Conclusory Implicational Roles in Implication-Space Semantics:

Recall from last time that a *premissory* role inclusion relation $A \subseteq_P B$ holds iff A can be substituted everywhere for B as the *premise* of an implication, *salva consequentia*, and a *conclusory* role inclusion relation $A \subseteq_C B$ holds iff B can be substituted everywhere for A as the *conclusion* of an implication, *salva consequentia*.

These metainferential relations are easy to express in the implication-space setting.

For $[A]=\langle a^+, a^- \rangle \in \mathbf{C}$ and $[B]=\langle b^+, b^- \rangle \in \mathbf{C}$,

$\mathbf{A \subseteq_P B}$ **iff** $\mathbf{a^+ \subseteq b^+}$ and $\mathbf{A \subseteq_C B}$ **iff** $\mathbf{b^- \subseteq a^-}$

In light of the observation last week that K3 is the logic of premissory role inclusions and LP is the logic of conclusory role inclusions, these facts indicate how to provide sound and complete semantics for those logics in the implication-space setting.

Common Structure is Commutative Monoids with a Distinguished Subset of the domain.

Monoids on a domain take any pair (or set) of elements into a further element of the domain.
Q: Why monoids in semantics?
For the general idea, think of semantically interpreting sentences on a *lattice* of them.
So every pair of elements has a meet and a join. Consquences is the lattice ordering, and that extends from atoms to their conjunctions (meets) and disjunctions (joins). But how does one introduce conditionals? Each conditional either holds in the prior lattice structure or it does not. We prematurely get conditional semantically evaluated, and have no representation of *false* conditionals. But we can introduce a conditional-forming monoid (think of it as adding elements *above the plane* of the original lattice. Some of those can be designated (perhaps those corresponding to the actual lattice ordering), but "false" conditionals can be represented, too.

<u>Metalinguistic Functionalism about Reason Relations:</u>

Four Universal Rational Metavocabularies

A. <u>Extrinsic-Explanatory</u>:

1. Pragmatic: Bilateral, Two-sorted Deontic Normative Pragmatic Metavocabulary
2. Semantic: Truthmaker Semantics with the Hlobil Consequence Relation

B. <u>Intrinsic-Explicative</u>:

3. Logical: NMMS codes an LX logical extension of arbitrary base vocabularies, specified in the metainferential vocabulary of the multisuccedent sequent calculus.
4. Conceptual Roles: Implication-Space Semantics.

Hlobil proves an isomorphism between (1) and (2) at the level of reason relations.
(4), Implication-space semantics characterizes what is in common between (1) and (2), and shows that that is also what is expressed by the logical vocabulary of (3).
Implication-space semantics is the vocabulary of pure conceptual roles or rational forms, which are articulated by reason relations.

Metafunctionalist Claim: Reason relations are just whatever can play these specific roles with respect to *all* four of these kinds of metavocabulary.

"We have developed two overarching and interrelated ideas.

The first of these ideas is that **propositional conceptual contents are articulated by open reason relations**, that is, by reason relations that needbe neither monotonic nor transitive. These reason relations are relations of implication and incompatibility, and they constrain the norms that govern practices of giving and asking for reasons. In virtue of standing to one another in reason relations of implication and incompatibility, things are bearers of conceptual contents—specifically *propositional* contents. Those contents can be thought of as the functional roles the sentences play in constellations of implications and incompatibilities. Declarative sentences express what can both serve as and stand in need of reasons (for and against, positive and negative). This is a version of **semantic inferentialism. Our semantic inferentialism includes a functionalism about conceptual contents: propositional conceptual contents are individuated by the role they play in reason relations.**

The second idea is that we can make explicit these reason relations in terms of different metavocabularies. In particular, **we can give accounts of reason relations in a pragmatic-normative metavocabulary and also in a semantic-representationalist metavocabulary**. We call these kinds of metavocabulary **"extrinsic"** because they appeal to conceptual resources that are not provided by reason relations themselves. There are also **two kinds of intrinsic metavocabularies**, which use only conceptual resources provided by reason relations themselves.

In particular, we can make reason relations explicit within the object language by using **logical vocabulary**, and we can also make them explicit in **the metavocabulary of implication space semantics**. The idea is that these four metavocabularies afford us four perspectives on a common topic, namely reason relations. We thus undertook to understand reason relations as what shows up from the perspective of all four metavocabularies. We called the roles played in reason relations that show up from all four perspectives "rational forms." [281-2]

a) Extrinsic-Explanatory vs. Intrinsic-Explicative Rational Metavocabularies.
   i. Implication Space Semantics is an *intrinsic* rational MV. It requires *only* the reason relations of the vocabularies for which it provides a semantics. The semantics is in that sense *internal* to the reason relations it uses to articulate conceptual roles of sentences (=claimables).
   ii. Sequent calculus means of extending base vocabularies to logically extended vocabularies also need as arguments *only* what is provided by the base vocabulary. Sequent calculi are pure (intrinsic) rational metavocabularies. The logical vocabularies they introduce are *sui generis*: explicative, but not their expressive functions as metavocabulary (explicating reason relations) are exercised in (an extension of) the object language. So they are a very special kind of MV.
   iii. Bilateral Deontic Pragmatic MVs appeal to doxastic commitments to accept/reject sentences of the lexicon of the vocabulary, and preclusions of entitlements from constellations of those commitments. But they aim to *explain* what the turnstile means, what implication (and reason relations generally) *is*, in terms of features of the *use* of expressions.
   iv. Truth-maker Semantics, appeals to states, fusion, and impossibility. MVs of the truth-maker sort aim to *explain* what implication *is* in terms of the impossibility of the states that result from fusing truthmakers of premises with falsity-makers of conclusions.

"The main differences between the theories lie in their explanations of what it means for contents to occur—in one of the two ways in which they can occur—and how and why some combinations of such occurrences are ruled out, where this ruling-out thereby exhibits different modal flavors.

For instance, (the right kind of) **semantic metavocabulary** explains the content expressed by non-logical base sentences by saying what features of the world those sentences represent. The sentence "The coin is made of pure copper" implies the sentence "The coin would melt at 1085 degrees Celsius" because and in the sense that it is alethically ruled out (it is impossible) that pure copper is solid at 1085 degrees Celsius, and it is incompatible with the sentence "The coin is an electrical insulator," because it is alethically ruled out for pure copper not to conduct electricity (pure copper necessarily conducts electricity).

Pragmatic metavocabularies explain what is expressed by non-logical base sentences by saying what features of the discursive practices of *using* those sentences it is in virtue of which practitioners count as practically taking or treating the sentences *as* standing to one another in relations of implication and incompatibility. The sentence "The coin is made of pure copper" implies the sentence "The coin would melt at 1085 degrees Celsius" because and in the sense
that it is normatively ruled out to accept the former but reject the latter (one cannot be entitled to that constellation of commitments), and it is incompatible with the sentence "The coin is an electrical insulator," because it is normatively ruled out to accept both sentences (one cannot be entitled to that constellation of commitments). The structure of these accounts is the same. What differs is the explanation of the relevant kind of ruling out of combinations, the relevant kind of exclusion between occurrences of contents.

If we view the alethic exclusion relations between truth-makers and falsity-makers and the normative exclusion relation between acceptances and rejections as two sides of one coin, we thereby also view the pragmatic normative and the semantic-representationalist theories of conceptual content as two sides of one coin. We can indeed view the alethic and the normative exclusion relations as two sides of one coin by holding that both pragmatic and semantic metavocabularies can be regarded as *rational* metavocabularies and, hence, as having at least one common topic, namely reason relations. Among the things they both discuss and seek to explain in their own terms are reason relations.

This way to reconcile the two traditions in the philosophy of language becomes compelling if we say that what matters for having the topic of reason relations in view is the structure that is shared between the pragmatic-normative and the semantic-representationalist theories. That is **metalinguistic functionalism about conceptual contents**. If and insofar as this line of thought is correct, we should be interested not only in the way in which reason relations show up both in the norms that govern acceptances and rejections and in the modal relations among worldly states, but we should also be interested in the structure of reason relations in abstraction from these different ways in which they can be enmattered." [285]

<u>Implication-space semantics and intrinsic rational MVs</u>:

"**The metavocabulary of implication-space semantics makes explicit reason relations in their unenmattered and pure form. In implication space semantics, we make explicit the roles that things play in reason relations, merely in terms of those reason relations themselves. The way in which we do this is by considering the ranges of subjunctive robustness of implications**, that is, the range of additions to the implication that yield or preserve a good implication. The relata of reason relations are treated as sets of implications in which only the given relatum occurs, as a premise or as a conclusion. The roles that relata of reason relations play are then defined as the equivalence classes of implications with respect to the equivalence relation of having the same range of subjunctive robustness— strictly, a pair of such roles, namely one in which the first member is the role of the implications in which only the given relatum occurs as a premise and the second is the corresponding role for the relatum occuring as a conclusion. In this way, we take reason relations and abstract from them the roles that relata play in them. We can then give an account of logical vocabulary just in terms of these roles. This account captures the common structure of the accounts of logical vocabulary given previously in semantic and pragmatic metavocabularies, and it does so while abstracting from any particular way in which reason relations are enmattered.

Implication-space semantics provides a powerful new perspective on reason relations and conceptual contents because it allows us to move from particular bearers of roles in reason relations to those roles themselves. We can define operations on those roles, such as adjunction and symjunction, and we can define new roles in terms of old ones. These roles abstract away from any particular matter, shape, structure, or syntax that their bearers may have. For instance, two sentences that differ markedly in their syntactic structure can play the same implicational role. So can sentences in different languages, vocabularies with different sets of sentences. (In these respects, **the conceptual contents of sentences as we understand them correspond to the functional classifications Wilfrid Sellars expresses with dot-quoted sentences such as •Der Tisch ist kühl• in his original inferentialist semantics**.)

We can interpret the pragmatic-normative and the semantic representationalist theories in implication-space semantics, and doing so reveals that corresponding sentences and worldly propositions in the two theories can play the same implicational roles. That is, corresponding sentences and worldly propositions play the same role in their respective reason relations. They accordingly have the same conceptual content.

Implication-space semantics reveals, in this way, the common structure of our pragmatic-normative and the semantic-representationalist metavocabularies. We take this common structure to underlie the representation relation between sentences and worldly propositions. The idea is that, for instance, the sentence "It is raining," which can occur in acceptances and rejections, represents the worldly proposition of it raining, which can occur in worldly states that make "It is raining" true and in worldly states that make it false, because the sentence and the worldly proposition have the same conceptual content—that is, they have the same implicational role.

To sum up, **the metavocabulary of implication-space semantics allows us to characterize reason relations in a way that brings out the common structure of explanations of reason relations in semantic and pragmatic metavocabularies**. In particular, we can now understand the pragmatic normative theory as a theory of which discursive acts users of concepts must treat as implying each other or being incompatible with each other in order thereby to treat their sentences as representing a particular worldly proposition.

Specifically, the exclusion relations among discursive acts must be isomorphic to the exclusion relations among worldly states, in spite of the different kinds of modality that articulate the reason relations of the two kinds. That is, combining acceptances or rejections of the sentence with other discursive acts and combining truth-makers or falsity-makers of the worldly proposition with other states must always be such that either both combinations are ruled out or neither is.

We thus have three different perspectives on reason relations: one in terms of the norms governing thought and talk, one in terms of the possibility or impossibility of worldly states, and one that characterizes the roles that things play in reason relations abstractly. The third perspective binds the first two together, making them visible as providing different perspectives on one topic: reason relations.

Logical vocabulary provides the fourth perspective. Indeed, it has guided our path in binding the first two perspectives together in the third one. For **it is the treatment of logical vocabulary that brings out in detail how the three perspectives relate to each other**. Once we see how the rules for, say, the conditional or conjunction as stated in our three metavocabularies correspond to each other, the overall correspondence between the theories formulated in these metavocabularies becomes clear. This is so because logical vocabulary makes explicit the reason relations of arbitrary base vocabularies. As a result, by finding the correspondences between the theories of logical vocabulary in our three metavocabularies, we can appreciate the correspondence between the reason relations for the entire language.

Logical vocabulary is not a metavocabulary, in the sense that it does not talk *about* reason relations or its relata. Rather, it makes reason relations explicit within (an extension of) the object language itself. The vocabulary of the sequent calculus, which we use to introduce logical vocabulary as universally LX, *is* a rational metavocabulary in the strict sense. Logical vocabulary is, nevertheless, a rational metavocabulary in the sense that it allows us to form sentences whose undeniability makes explicit what implies what and what is incompatible with what." [287-88].

b) Reason relations as what can be specified in all four of these fundamental ways.
   i. Q: On what assumptions are these a *complete set*? That is, could there be others, a third extrinsic-explanatory kind of rational MV? What about other intrinsic-explicative ones?
   ii. Do pragmatics and semantics exhaust the relevant space of extrinsic-explanatory rational MVs? Is there room for one that appeals to pleasingness to God?

On the intrinsic-explicative front, *surely* there can be other mathematical specifications of reason relations: category-theoretic, intensional type-theoretic, or homotopy type-theoretic.