November 30, 2022

Notes for Week 13

Intro:

The material I'll present today is different from what we have seen so far, in a number of ways. First, I've been surveying ideas that we have been able to make work, results we have been able to prove, or things we are more or less satisfied we have figured out.

Today I'll be talking about two ongoing projects, more or less promising, that are still *medias res*.

1) The first is a computer implementation of what we call "dialogic pragmatics."

Here we do have a result, in the form of an up-and-running program.
It does what we set out to do, and proves an important point.
But we are in the very early stages of learning what we can in principle learn from it about the interactions between formal features of abstract reason relations and different aspects of discursive practices conducted in accordance with norms articulated by those reason relations.

2) The second is a long-time fantasy of mine about the possibility of logics that are expressively complete codifications of reason relations in yet another sense of "expressive completeness"—in addition to Dan's representation theorem showing the relations between sets of sequents in logically extended languages and features of the underlying material reason relations that are logically extended.

This is the idea of what I call a "monadologic."
By that I mean a logic such that the logically complex consequences of *each and every* single premise-set encode the entire set of reason relations concerning *all* possible premise-sets in that vocabulary.

And here—spoiler alert!—I should emphasize that we do *not* know how to achieve this sort of expressive completeness, and do not even know whether the monadological expressive aspiration is actually coherent.
It might be that there are deep reasons why monadologics are impossible in principle.
But I do have some ideas about how it might be done, and at least one concrete proposal that has some desirable and promising consequences that I *can* demonstrate.
So I want to share with you the current state of play for this pie-in-the-sky project.

I. Dialogic Pragmatics

<u>Plan</u>:

1) Introduction
2) Modeling Vocabularies and Reason Relations
3) DP1
4) Using DP1 as a probe to study the internal structure of reason relations.
5) Aspirations: DP2 (and DP3)

1. <u>Introduction:</u>

Early on, I started to talk about "vocabularies."
These are not to be identified with their *lexicons*, the set of *sentences* they contain.
They are supposed to be sentential lexical items that *mean* something, because of the way they are *used*.

Later on, I offered a technical, *algebraic* sense of "vocabulary":
A vocabulary is a *lexicon* plus a set of *reason relations*.
The lexicon is a set of sentences, about which all that matters is how many of them there are, and that we can tell them apart. So they might as well be sentence letters ('A', 'A_4', 'p') or just numerals.
The reason relations are a set of *implications* and a set of *incompatibilities*.
The only structure we impose is to require that the incompatibilities be *symmetric*.
One way to ensure that is to derive them from a set of *incoherent* sets, by the principle that two sets are incompatible with each other, iff their union is incoherent.
(We have seen that we *can* encode incompatibilities in the implication relation, by using empty RHSs. But this trick is entirely optional.)
This technical notion of vocabulary amounts to identifying them as *relational structures* in the algebraic sense in which Tarskian model theory uses relational structures as models.
In this case, the domain is the lexicon, and the relations are (in the multisuccedent case) sets of pairs of sets of lexical items—in the intended interpretation, the good implications.

I also offered a *pragmatic* account of reason relations of implication and incompatibility bilaterally, in terms of practical attitudes of acceptance/rejection, and speech acts manifesting them of assertion/denial, along with two kinds of deontic normative statuses of *commitment* and *preclusion of entitlement*.

This last contrasts with the single-sorted deontic status of in-bounds/out-of-bounds that Restall/Ripley bilateralism uses.

I also sketched an account of minimal necessary conditions on a social practice counting as a *discursive* practice.
Some speech acts must have the significance of *claimings*, which we can think of as assertions or denials, expressed by speech acts of producing sentences in one of two modes (asserting/denying).
It must be possible to rationally *challenge* claims, and to rationally *defend* them.
Challenging claims is making further claims that offer reasons *against* the challenged claim.
Defending claims is making further claims that offer reasons *for* the defended claim.
Reasons for accepting a claim are sentences that *imply* them.
Reasons against accepting a claim are sentences that are *incompatible* with them.
I suggested that a default-and-challenge structure of *entitlement* to *commitments* to accept and reject is the necessary basis for discursive practices.

We have seen how to go from discursive practice of making claims (assertions/denials) and defending and challenging them rationally by giving reasons for and against them to vocabularies, as sets of reason relations of implication and incompatibility defined on a lexicon.
This is our (Simonelli-inspired) version of RR-bilateralism, with the additional fine structure of a two-sorted deontic normative metavocabulary.
A premise-set Γ:

- *implies* a sentence A iff commitment to accept everything in Γ precludes entitlement to reject A, and
- is *incompatible* with A iff commitment to accept everything in Γ precludes entitlement to accept A.

So we have both have defined "vocabularies" algebraically as a lexicon plus a set of reason relations, and also
talked about vocabularies as sitting between "languages" and "theories" as Rorty thinks we must do after Quine's pragmatic undercutting of the language/theory, meaning/belief distinctions.
The point is seconded by Wittgenstein, seeing agreement in judgments as an essential element of agreement in meanings.

Here the point is that vocabularies are not just lexicons.
They are meant to be lexicons *in use*, or *as used*.
For underlying the Wittgenstein-Quine point is the idea that meanings are manifested only in actual practices of making claims and defending and challenging them.
But it seems that I have given a very abstract, technical sense to "vocabulary."
True, vocabularies are not just lexicons.

But adding reason relations, for instance, as sets of pairs of sets of subsets of the lexicon, does not get us very close to the *use* of the vocabulary.

So it is sensible to ask:
Given that we can go from discursive practices to vocabularies, can we go the other way around?
Does specifying a vocabulary in my spare, technical algebraic sense (intimately connected to Dan's implication-space semantics), actually set norms for discursive practices of asserting/denying, challenging/defending, altering the "score" of commitments and entitlements to attitudes of acceptance/rejection.?

If the story all hangs together properly, then specifying a *vocabulary* in the abstract mathematical sense should determine norms governing discursive practices

But does it?

It seemed to me that the most cogent justification of the claim to *pragmatic* adequacy of the connection we have put in place between reason relations and discursive practice would be to show that the mere capacity to associate reason relations with lexical items, that is, to specify a vocabulary, can be *algorithmically elaborated* into the capacity to engage in discursive practices meeting the conditions we have laid down.

The best way to show the possibility of algorithmic elaboration is to write a computer program that does it.

So we did.
Well, Pitt philosophy Ph.D. student Yao Fan did, under my supervision—and Pitt philosophy Ph.D. student Dan Webber now maintains and extends the program and its descendants.
(I actually wrote some bits of code, but even the bits that did what they were supposed to do were so clumsy that Yao rewrote them all in his elegant, efficient code.)
The program I will demonstrate is our proof of concept.

The Python program DP1 (for "Dialogic Pragmatics 1") takes as input a vocabulary in my technical, algebraic relational structure sense of lexicon plus implications and incompatibilities (reason relations), and conducts a dialogue that consists of speech acts of asserting/denying, with the pragmatic significance of challenging, and defending, and the keeping of deontic score on how those speech acts institute and alter deontic statuses of two kinds, commitment and entitlements, relative to bilateral attitudes of acceptance/rejection.
In fact, we specify a lexicon (a set of numerals), and the program generates reason relations of implication and incompatibility among them, randomly within the constraints we give it.

The program then generates dialogues in which claims are made, challenged, and defended, and deontic scores altered accordingly.

2. <u>Modeling Reason Relations (Vocabularies):</u>

We have so far considered reason relations in a very abstract way.
The lexicon of a base vocabulary is a finite set of "sentences" that are thought of as merely numerically distinct—so representable by sentence letters, or, if there are not enough of them, as numerals.
This is good enough, because the so-called "sentences" are not thought of as *meaning* anything, in advance of consideration of the role they play in reason relations of implication and incompatibility.

In the implication-space semantics, the reason relations defining a *vocabulary* on that lexicon is then the set of all pairs of subsets of the lexicon—thought of as *candidate implications* in a multisuccedent system.
A subset of those candidate implications is then distinguished as the *good* ones, the implications that actually hold. Those that have empty right-hand sides mark their premise-sets as *incoherent*, and thereby determine the material incompatibilities of their subsets.

But we never looked at any *actual* lexicons or reason relations.
In particular, we never looked at their internal structure, except to argue that we should *not* impose global closure structural principles of monotonicity and transitivity on the reason relations of material base vocabularies.

Now we want to look more closely at specific sets of reason relations.

We do that, for demonstration purposes, with a lexicon of just 7 sentences, in the form of the numerals '0' through '6'.

What you see on your handout is a set of implications and incompatibilities defined on that lexicon: an "actual" set of reason relations—or, if that seems an odd description of relations on a bunch of numerals, a *concrete model* of a set of reason relations on a lexicon.
That is, a *vocabulary*.
This one is computer generated, not imposing closure structure, and omitting implications that are good because of CO—that is, because the conclusion is contained in the premises.
(For simplicity, we do all of this with single-succedent reason relations.
This involves no restriction in principle, since we know how to generalize from this case.)

I want to pause a bit to talk about such concrete models of reason relations, because there are some lessons to be learned already just from thinking about them.

a) First of all, we should appreciate just **how many vocabularies** are definable on this very small, spare lexicon. Seven sentences is not a lot. Even a simple English sub-grammar working on the 5,000 word dictionary of Basic English will produce many millions of sentences of less than, say, 10 words.

But our poor 7-membered lexicon has 2^7, which is to say, 128, subsets.
To model incoherent sets, we pick a set of such subsets, to be paired with the empty set as a conclusion.

How many choices are there? How many sets of subsets of the lexicon (subsets of its powerset) are there? 2^{14} is more than 10 million—10^7—(since 2^{10} is 1024, approximately 10^3). And we have to make such a choice—one of the 10 million sets of subsets of the lexicon—to be premise-sets paired with each of the 7 single sentence conclusions in the lexicon.
These choices are, so far as anything we have said, independent of one another, so to get the total number of possible vocabularies, we need to multiply these together, yielding 10 million to the 8th power. 10^7 to the 8th power is 10^{15}. That is a *lot*.
So even our impoverished 7-element lexicon generates a very large number of possible vocabularies.

b) And at this point we are in a position to be reminded of our stunning ignorance about *actual* reason relations in natural languages. For the next decision in building a *concrete model* of reason relations for our toy vocabulary is to decide how *many good* implications and incompatibilities there should be in a vocabulary with a lexicon of n sentences.

Are one out of ten possible implications good—in, say, nautical vocabulary, or culinary vocabulary, or theological, or geological vocabularies? One out of a hundred? A thousand?
And whatever that number is, how is it related to the number of incompatibilities?
Should we expect more reasons *for* (good implications) than reasons *against* (incompatibilities)?
What is even a *range* of sensible ratios between them, and proportions of good to possible reason relations more generally?
At a finer grain, looking sentence by sentence, if we fix a conclusion, what range of ratios of possible reasons *for* that conclusion to reasons *against* that conclusion are sensible.
What difference does it make if there are wide differences between different sentences in this respect?

There is a public database at Stanford of ¾ of a million pairs of English sentences, about which Task-Rabbit testers were asked whether one implied the other, or was incompatible with the other, or neither. But their sentences involve many different topics, and they made no attempt to mix and match them so as to make possible an estimate of how many of the *possible* combinations would have been endorsed as implications, incompatibilities, or irrelevancies.

So we are working in the dark here, and can only experiment with different values of these important *ratios* of reason relations—of reasons for and against, and of good reasons as opposed to irrelevancies.

3. DP 1:

What the program DP1, which I am about to demonstrate, does is take a set of reason relations of the kind of which you have a sample, and generate dialogues in which interlocutors make claims (assertions or denials of lexical items) and give reasons for and against them, so defending and challenging those claims.
The program keeps track of commitments undertaken to accept or reject claimables, and implements a default-and-challenge structure of entitlement to those commitments depending upon the success of rational challenges and defenses of those commitments.
The initial reason relations provide the stock of implications and incompatibilities—treated as common to both interlocutors—from which reasons for and against can be drawn to defend and challenge commitments undertaken, by undertaking further commitments.

We can now look at the sample summary of a dialogue that is on the handout, which will help in understanding what the program produces.

Then I'll walk through one such dialogue step by step, and then show how we can re-run such a dialogue from any given point, to see what happens if different moves were made at various points.

[Do that.]

4. Some lessons and questions:

a) The first conclusion I want to draw is the most important: it *is* possible to derive appropriate norms governing social discursive practices of making claims and giving and asking for reasons for and against them, rationally defending and challenging those claims, and keeping track of the evolution of commitments and entitlements over the course of such moves in a conversational game.

(Dan Webber has observed that this one is modeling something like a **high-school debate**.)
So we have made available conceptual resources sufficient to go in both directions:

- to understand abstract reason relations of implication and incompatibility in terms of discursive practices of making claims and rationally defending and challenging them by offering reasons for and against them, and

- to understand how abstractly specified reason relations (vocabularies) provide norms governing such discursive practices.

b) I pointed out earlier that we are stunningly ignorant about very basic facts about actual reason relations of material implication and incompatibility. One aspiration of DP1 is to use the course and results of various dialogues as a *probe* to illuminate significant features of different vocabularies (sets of reason relations) defined on a common lexicon. (Think of analyzing the internal structure of subatomic particles by experimentally investigating the results of colliding them with others.)
What respects of structural similarity and difference among reason relations correspond to *pragmatically significant* differences in dialogues conducted according to the norms of reasoning they codify?

Here it is a big advantage that we can run many thousands of dialogues using the same set of reason relations, and see what regularities emerge from them.
We can then compare those to the regularities that emerge in dialogues governed by reason relations that differ from those in specified ways.

The enterprise of studying these relations—and so of finding out what structural features of vocabularies have various sorts of pragmatic significances for discursive practice—is in its infancy. We have barely begun.
The tentative experimental observations I'll report were made a couple of years ago, and we have not actively revisited this part of our project since. Life is full.

It seems that there are large differences between different sets of reason relations (vocabularies on a common lexicon) with respect to at least two of the output variables concerning dialogues that DP1 runs, namely success at defending commitment to accept or reject specific claimables, and common ground between the interlocutors that emerges from a dialogue. By "common ground" here I mean shared commitments to accept or reject claimables, with durable entitlements. While contesting a particular commitment, interlocutors almost always end up agreeing to accept or reject certain *other* claimables. We can track the relations between these common grounds and the original commitment, and see how these vary across different sets of reason relations.

i. Some sets of reason relations make it *much* easier to sustain entitlement to *some* commitments than others.
ii. Some vocabularies make it easier for different interlocutors to *agree* about and become jointly entitled to some claims, rather than others.

In general, this is probably not surprising. After all, some vocabularies incorporate many more reasons for a given claim than they do reasons against it, or *vice versa*.
But as far as we can tell, these differences in what we could call "reason ratios" do not nearly account for the differences in sustainability and common ground that we observe.

Those differences must depend on subtler interactions of reasons for and against one claim with reasons for and against others.
It is exactly those interactions we would like to understand better.

As I put the point a minute ago, we would like to find out what respects of structural similarity and difference among reason relations correspond to *pragmatically significant* differences in dialogues conducted according to the norms of reasoning they codify?
As things stand, we don't even have a good set of *candidate* formal properties of reason relations to consider.
The hope is that experimenting with different vocabularies that yield substantial differences in pragmatic effects will give us clues as to what descriptions of those reason relations pick out *important* respects of similarity and difference.

The following train of thought might be fanciful, but I see a connection here with traditional philosophical concerns with what have been called "**coherence theories of truth**."
The vague thought behind such "theories" is that what it is for a claim to be true is for it to hang together rationally with other claims in specific ways.

Although it cannot be said, I think, that anyone ever actually *worked out* such a theory, the idea is that the truth-talk is a misleading way of talking about how the elements of a rationally *coherent* set of commitments "hang together" by providing reasons for each other, and excluding commitments to which one is *not* entitled by the privileged ones.

One way of following out that idea is to think that *if* one understood all the reason relations in which sentences stand to one another, *then* one would understand with ones one could be entitled to be committed to accept and reject. Being *justified* in that holistic sense is then offered as a notion that can do the explanatory work that the concept of <u>truth</u> (understood, say, as a kind of correspondence to something else) has traditionally been called on to do.
Thought of this way, the project of a coherence theory of *truth* has two parts:

- First, offer a holistic coherence theory of *meaning* or content, in terms of role in the sort of reason relations that matter for *justification*.
- Then, understand truth of *sets* of claimables in terms of ideal justificatory relations that they can stand in to one another.

Probably the sensible view here is that the first move is worth trying, but the second move is a step too far.

Speaking against that sensible view, though, is the observation common to Wittgenstein and Quine that agreement in the meanings or contents of claimables is not wholly separable from agreement in commitments to constellations of them.

I am suggesting that against this background, one might find it interesting to add the datum—which we are in a position to demonstrate and to explore in detail—that some sets of reason relations make it *much* easier to justify and sustain entitlement to, and to agree in one's verdict on, some sentences rather than others.
This is not an all-or-none, true or false difference.
But the connection it forges between meaning, justification, and accessible agreement is both suggestive and one we can investigate in hitherto undreamed-of detail and precision—even within our ludicrously simplified setting.

5. Aspirations from Girard's Ludics

DP1 is based on a clear picture:

II. Discursive practitioners agree entirely on the reason relations that govern the sentences they use.
III. They engage in practices of making claims and rationally defending and challenging them by giving reasons for and against their commitments to explore which commitments one can sustain entitlements to, in the light of critical attention.

(Side note: DP1 does allow that the interlocutors might have different "inferential theories," in that they each accept only *part* of the whole set of reason relations—parts which overlap, but need not wholly coincide. But we have not yet experimented much with this feature.)
This is a movement, roughly, from meaning to belief.
But I buy into the Quinean point that natural languages don't obey the two-stage procedure Carnap described for artificial languages: first specify the meanings and then determine which sentences with those meanings turn out to be true.
All we do is *use* sentences, *both* making claims and rationally defending and challenging them, appealing to reason relations that articulate their meanings.
That thought suggests that a more realistic simulation of discursive practice would have the reason relations emerge gradually through dialogue, developing and filling-in in tandem with commitments and entitlements to claimables by various interlocutors.

Programs DP2 and the nascent DP3 (which Dan Webber is working on) are attempts to implement that idea.

I'll just show a bit of how DP2 works.
Interlocutors start with some commitments they would like to be able to sustain, and different strategies for endorsing implications and incompatibilities—incorporating them into the growing set of agreed-upon reason relations.
There is a summary of this on the handout.

[Brief demonstration of DP2.]

The feature of DP2 that is of the most interest, perhaps, is that we incorporated a distinctive kind of *negotiation* over acceptance of particular reason relations.
This is illustrated by the “tree” structure.
The idea is that one interlocutor might say:
“Would you agree that ‘I strike this dry, well-made match’ implies ‘The match will light.’?”
And the other might respond:
“I agree to that implication if you agree that ‘I strike this dry, well-made match’ and ‘There is a very strong magnetic field around me and the match,’ does *not* imply (or even is incompatible with) ‘The match will light.’”
To which the original speaker might respond with a further offer of conditional acceptance:
“I will accept both those reason relations if you will agree that if, in addition, I am in a Faraday cage, then it *would* follow that the match would light.

DP3 is trying out quite a different way of generating sets of reason relations by dialogical interaction, relying on metainferential patterns.
Really early days on this one.

It is perhaps worth mentioning that we have successfully daisy-chained DP1 and DP2, by using the reason relations that result from dialogues in DP2 as inputs to the dialogues in DP1.
We think combinations of this kind might be much better models of the phenomenon Wittgenstein and Quine point out about the interaction of meaning and belief than either sort of program is on its own.
But, again, we’ve so far only dipped our toes into this aspect of the project.

IV. Monadologics

Start with firm disclaimer:
Unlike the other stuff I've talked about this term, I do *not* know how to do what I am talking about doing here. And in that significant sense, I don't even know whether what I am talking about is ultimately *coherent*. It could well be that in the end it turns out to make no sense.
But if it *does* make sense (a very peculiar sort of subjunctive!), then there is what seems to me to be an exciting project here. So I am going to talk about it as though it *did* (does) make sense, until we understand things better and can see why it does not.

Plan:
a) Two views of the intensionality of our conditional and negation
b) The idea of monadologics and
c) The means: the downward conditional.
d) Criteria of adequacy for downward conditional and negation.
e) Significance of monadologics for database management: turning online updating into offline updating.

a) two views of the intensionality of our conditional and negation

The principal sentential connectives of the logical system NM-MS (the conditional and negation: the expressive, rather than the merely aggregative connectives) are semantically *intensional* in that they make what follows from *one* premise-set depend on what follows from some *other* premise-set. Here we are thinking of premise-sets as points of semantic evaluation (think: worlds, though premise-sets correspond at best to *partial* worlds). The *logically complex* consequences of, what is *implicit* in, a premise-set, depends on what is implicit in *other* **nearby** premise-sets.
So, paradigmatically, whether or not Γ |~ A→B depends on whether or not Γ,A |~ B.
Whether or not Γ |~ ¬A depends on whether or not Γ,A |~ , that is, whether or not the *different*, but adjacent premise-set Γ,A is incoherent.
In both cases, to tell what logically complex sentences follow from Γ, one must look at what follows from its neightbor premise-set $\Gamma \cup \{A\}$.

We can turn this observation around, and say that the point of introducing logical vocabulary is to encode, in the consequences of *one* premise-set, information about the *suppositional neighbors* of that premise-set.

In fact, as we can see from the remarks about conditionals and negations above, for any sentence X that one wanted to add to Γ, the logically complex consequences of Γ in the form of

conditionals with X as antecedent and negations with X as negated (in both cases, as principal connective) tell us about the consequences and coherence of premise-set Γ,X.
We can repeat that process, adding Y, and then Z. The logically complex consequences of Γ will then encode information about the consequences of Γ,X,Y, and Γ,X,Y,Z, and so on.
In short, the logically complex consequences of Γ encode the consequences of *all* of Γs supersets—right up to the whole language L (which we stipulate is incoherent, and which by CO implies everything).
Further, Γ does not just codify or encode *some* information about the consequences of all the premise-sets that are its supersets. It codifies *all* the information about those consequences.
If $\Gamma,\Delta \mid\!\sim A$, then $\Gamma \mid\!\sim \&\Delta \rightarrow A$ and
If $\Gamma,\Delta \mid\!\sim$, then $\Gamma \mid\!\sim \neg(\&\Delta)$,
(where &Δ is the conjunction of all the elements of Δ).
(We assume throughout that all premise-sets are finite, though their size is otherwise unbounded.)

This is the *expressive point* of *intensionality*.
The inferential role of sentences formed using intensional sentential operators expresses information about the inferential roles of the component sentences to which those operators are applied.

These operators (conditional and negation, but also the monotonicity modality box) let us *say* how things are in the neighboring regions of the space of suppositions: premise-sets from which to explore and extract consequences.

I have just argued that our conditional and negation (in virtue of their right rules satisfying dual Ramsey, that is, Deduction theorem and Detachment (DD), and Incoherence-Incompatibility (II)—see Week 6 Handout) are *expressively complete* codifications of the consequences-and-incompatibilities (reason relations) of the *supersets* of the premise-sets that imply logically complex sentences formed using them.

b) **The idea of a monadologic**:

The two observations—

i) about the expressive role of these intensional sentential operators and

ii) about its being expressively complete only for supersets of each anchor premise-set—

invite a further expressive ambition.

For these logical operators do *not* permit expression, in the logically complex consequences of a premise-set Γ, of the reason relations of Γs *sub*sets, nor of premise-sets *disjoint* from Γ, or merely *overlapping* Γ.
They are exclusively *upward*-looking (it's face raised to heaven), codifying (all) of the reason relations *only* of Γs supersets: its *upward cone* in the lattice of subsets of the language.

In *this* sense, our conditional and negation (with their aggregative helper-monkeys of conjunction and disjunction) are an expressively *in*complete *set* of logical connectives.
For they only codify in the consequences of *one* premise-set the reason relations in which *some*, but *not all* other premise-sets stand.
An *expressively complete* ***set*** of logical connectives would do for *every* possible premise-set defined on the lexicon of the material vocabulary in question what *our* connectives do for the *supersets* of a given premise-set: codify the reason relations of *all* premise-sets in the logically complex consequences of *each* premise-set.
This is quite a different sense of "expressive completeness" than the one Dan Kaplan proves for NM-MS relative to base vocabularies

This is what I call a "monadologic."
I derive the term from Leibniz's monads.
Recall that every one of his monads reflects its entire universe.
Their differences consist in *how* they do that: the degree of confusion or inadequacy of each perception and the variety of such degrees across the whole spectrum of a monad's perceptions.
But, taken together, all of each monad's perceptions represent the whole world, in all its variety.
I take this to mean that it is in principle possible to infer, from the class of all the perceptions of any single monad, what the perceptions of all the other monads is.

A ***monadologic*** would be a set of connectives such that the logically complex consequences of *each and every* premise-set codify the implications and incompatibilities-incoherences of *every other* premise-set, in the same sense in which the conditional and negation of NM-MS do for the *supersets* of the premise-set whose logically complex consequences they articulate.

In case you don't like the Leibnizian conceit, here is another metaphor for what I am after.
A monadologic would be *completely* and *perfectly holographic*.
Here the contrast in question is between *standard pictorial* transformations of a *scene* into an *image*, and *holographic* transformations of a scene into an image.
In the standard pictorial case, if you remove a contiguous portion of the representation, say cutting 10% of it off at the corner of the film, then you lose 100% of the corresponding 10% of the image. You can't see any of the image corresponding to the part of the film you snipped off.

But in a holographic transformation, removing any 10% of the film (whether contiguous or not) just removes 10% of the information from the whole thing. It just gets 10% fuzzier, has 10% less resolution.
(The original and still in some sense standard way of doing this is with **Fourier transforms** of arbitrary waveforms.)
In this case, *all* the information about the reason relations of the whole vocabulary is contained in the logically complex consequences of *any* and *every* premise-set.
Each by itself carries all the information needed to reconstruct all the reason relations of the entire vocabulary.
That is what I mean by saying that a monadologically extended vocabulary would be *completely* and *perfectly* holographic: as long as the logically complex consequences of *any* premise-set remain, one can remove *all* the rest without losing *any* information that the whole vocabulary contained to begin with.

c) **the means: the downward conditional.**

Let me emphasize that since we haven't built one—we don't have an up-and-running system with these expressive powers—I don't know whether there can be a monadologic in the sense in which I am sketching.

Here is the idea I have suggested (which I don't know whether can be made to work):
I observed above that both our expressive connectives "look upward", in that logically complex sentences formed from them are implied by Γ just in case supersets of Γ have certain features (implications and incoherences).
What we want, then, is corresponding connectives that "look downward."
That is, instead of codifying the effects of *adding* a premise A *to* Γ (get B as a consequence, or get an incoherent premise-set) we want to codify the effects of *subtracting* a premise A *from* Γ.
Recall that, because we are working in *open*-structured reason relations, *subtracting* a premise from Γ can result in *new* implications, implications of sentences *not* implied by the superset Γ of Γ-A. And it can be that while Γ is coherent, Γ-A is *in*coherent. (In that case, A is a *defeater* of the incoherence of Γ-A, which turns out not to be *persistently* incoherent.)
But we would also

The "upward-looking" character of our conditional is codified in the Deduction-Detachment (DD) (Dual Ramsey) criterion of adequacy on the right rule for the conditional:
Deduction-Detachment (DD) Condition on Conditionals: $\Gamma |\!\sim A \rightarrow B$ iff $\Gamma, A |\!\sim B$.
We want a *downward conditional* satisfying something like : $\Gamma |\!\sim A \rightarrow^{-} B$ iff $\Gamma\text{-}A |\!\sim B$.

As with NM-MS, we would codify these with canonical simplifying rules, which we would like to be reversible, Ketonen-style:

R$\rightarrow^{-}$: $\frac{\Gamma\text{-A} \mid\!\sim \Theta,\text{B}}{\Gamma \mid\!\sim \Theta,\ \text{A}\rightarrow^{-}\text{B}.}$

We need to decide what to do about the downward connectives in case A∉Γ.
In the case of the upward connectives, we decided that if A is already in Γ, Γ implies A→B whenever Γ implies B.
By parity, then if A is not in Γ, we would have Γ imply A$\rightarrow^{-}$B in case Γ implies B.
That is treating Γ-A as just being Γ when A∉Γ.
(Alternatively—and we have sometimes found this technically helpful in past tries—we can rewrite the LHS of the top lines of these rules as Γ,A -A. If A∈Γ, then Γ,A is set-theoretically just Γ, so long as we aren't working with multisets. So in that case Γ,A -A is equivalent to Γ-A. And if A∉Γ, then Γ,A -A is equivalent to Γ, which we want to imply A$\rightarrow^{-}$B just in case Γ already implies B, by the convention above.)

Suppose this works. (More on what "works" means here coming below.)
Then these "downward" connectives would allow us to capture, in the implications by Γ of logically complex sentences, the implications and incoherences of *subsets* of Γ, in a way parallel to how the "upward" connectives of NM-MS codify the implications and incoherences of *supersets* of Γ in its implications of logically complex sentences.

But functioning downward conditionals would be able to express much more than the reason relations only of subsets of Γ.
Combined with the upward connectives, the downward connectives would make it possible to express in the consequences of each premise-set Γ the reason relations of *all* premise-sets: not just subsets and supersets of Γ, but also the reason relations of premise-sets that only overlap Γ, and even those that are disjoint from Γ (in their explicit content, that is the members of the premise-sets).
For any Γ,Δ∈L, Δ is *accessible from* Γ by a series of *additions* to Γ of elements of Δ that are not in Γ, and *subtractions* from Γ of elements of Γ that are not in Δ.
Since we require premise-sets to be finite (though otherwise unbounded in size), these accessibility paths will also be finite, and so codifiable in finite conjunctions of sentences in the antecedents of conditionals and the prejacents of negations.
(If it became important to us to relax the requirement of finitude of premise-sets, our upward-and-downward monadologic would have each premise-set codifying the reason relations only of all *merely finitely different* premise-sets. Suitable compactness results might well make this restriction amount to no restriction at all. Here a property we care about is *compact* if its holding for all finite subsets of a set suffices for it to hold for the infinite set of which they are subsets.)

Note that for this step-up plus step-down procedure to yield unique results, we would have to prove a very strong *path-independence* result. That is, that if $\Delta = \Gamma + A - B$, that
$\Gamma \mid\sim A \rightarrow (B \rightarrow^{-} C)$ iff $\Gamma \mid\sim B \rightarrow^{-} (A \rightarrow C)$. For only then will the result of stepping up by A and then down by B be the same as the result of stepping down by B and then up by A.
In this respect, it is encouraging that in NM-MS, with just the upward conditional,
$\Gamma \mid\sim A \rightarrow (B \rightarrow C)$ iff $\Gamma \mid\sim B \rightarrow (A \rightarrow C)$ iff $\Gamma \mid\sim (A \& B) \rightarrow C$.
[I show this below.]

d) Criteria of adequacy for downward conditional and negation.

What will it take to make this work—that is, to implement a true monadologic by adding downward conditionals and negations to NM-MS?

i. I have already mentioned that we will need to be able to prove a very strong path-independence result, showing that we get the same results no matter what order we apply upward and downward connective in.

ii. We will need to show that adding the downward connective rules to NM-MS produces a *conservative extension* of the base vocabulary, in that no new sequents (implications or incompatibilities) involving only sentences from the lexicon L_0 of the prelogical base vocabulary are licensed by the new rules. This is easy to ensure, since the rules (though we want to keep them reversible) do not include any *simplifying* rules. The sequents below the line always include new vocabulary that the sequents above the line do not.

iii. We need to show *consistency* of the resulting system. Note that in a structurally closed setting, showing the conservativeness of the extension in (ii) would be enough for this. For it requires that some sequents *not* be provable—namely any involving only nonlogical lexical items that are not sequents ratified by the original base vocabulary. But in our structurally open (nonmonotonic, nontransitive) setting, this does not suffice to show that the resulting system is consistent, in that no sequent can both be shown to hold and shown not to hold.

iv. Perhaps most demandingly, we need to show that CO, Containment, holds, that is, that for every Γ:
$\Gamma, A \rightarrow^{-} B \mid\sim A \rightarrow^{-} B$ and $\Gamma, \neg^{-} A \mid\sim \neg^{-} A$.

Particularly this last requirement turns on the *left* rules for the downward connectives.

Our motivations extended only to the *right* rules—as indeed was the case with the original upward conditional and negation, with expressive criteria of adequacy DD and II.

Here, for reference is NM-MS:

L$\rightarrow$: $\dfrac{\Gamma \mid\!\sim \Theta, A \qquad B, \Gamma \mid\!\sim \Theta}{A \rightarrow B, \Gamma \mid\!\sim \Theta}$ R$\rightarrow$: $\dfrac{A, \Gamma \mid\!\sim \Theta, B}{\Gamma \mid\!\sim \Theta, A \rightarrow B}$

L&: $\dfrac{\Gamma, A, B \mid\!\sim \Theta}{\Gamma, A \& B \mid\!\sim \Theta}$ R&: $\dfrac{\Gamma \mid\!\sim A, \Theta \qquad \Gamma \mid\!\sim B, \Theta}{\Gamma \mid\!\sim A \& B, \Theta}$

L$\vee$: $\dfrac{A, \Gamma \mid\!\sim \Theta \qquad B, \Gamma \mid\!\sim \Theta}{A \vee B, \Gamma \mid\!\sim \Theta}$ R$\vee$: $\dfrac{\Gamma \mid\!\sim A, B, \Theta}{\Gamma \mid\!\sim A \vee B, \Theta}$

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

Let us add the downward conditional $\rightarrow^{-}$ with the two new rules:

R$\rightarrow^{-}$: $\dfrac{\boldsymbol{\Gamma\text{-}A \mid\!\sim \Theta, B}}{\boldsymbol{\Gamma \mid\!\sim \Theta, A \rightarrow^{-} B.}}$

L$\rightarrow^{-}$: $\dfrac{\boldsymbol{B, \Gamma\text{-}A \mid\!\sim \Theta}}{\boldsymbol{A \rightarrow^{-} B, \Gamma\text{-}A \mid\!\sim \Theta.}}$

<u>CO Preservation:</u>

$$\dfrac{\dfrac{\boldsymbol{B, \Gamma\text{-}A \mid\!\sim B, \Pi} \quad \textbf{CO}}{\boldsymbol{A \rightarrow^{-} B, \Gamma\text{-}A \mid\!\sim B, \Pi} \quad \textbf{L}\boldsymbol{\rightarrow^{-}}}}{\boldsymbol{A \rightarrow^{-} B, \Gamma \mid\!\sim A \rightarrow^{-} B, \Pi} \quad \textbf{R}\boldsymbol{\rightarrow^{-}}}$$

This proves CO for $\rightarrow^{-}$.

And it holds whether or not $A \in \Gamma$.

Here we can notice that along the way we in effect proved detachment from downward conditionals (DDC):

Γ-A, A$\rightarrow^{-}$B $\mid\!\sim$ B.

$$\dfrac{B, \Gamma\text{-}A \mid\!\sim B \quad \text{CO}}{A \rightarrow^{-} B, \Gamma\text{-}A \mid\!\sim B \quad L^{\rightarrow -}}$$

Since we have CO preservation for the downward conditional ((iv) above), the big outstanding result needed is path-independence of alternating upward and downward conditionals. My formulation above was:

Note that for this step-up plus step-down procedure to yield unique results, we would have to prove a very strong *path-independence* result. That is, that if $\Delta = \Gamma + A - B$, that
$\Gamma \mid\sim A\to(B\to^{-}C)$ iff $\Gamma \mid\sim B\to^{-}(A\to C)$. For only then will the result of stepping up by A and then down by B be the same as the result of stepping down by B and then up by A.
In this respect, it is encouraging that in NM-MS, with just the upward conditional,
$\Gamma \mid\sim A\to(B\to C)$ iff $\Gamma \mid\sim B\to(A\to C)$ iff $\Gamma \mid\sim (A\&B)\to C$.

Should show the analogue for $\to^{-}$ of the order-independence result for iterating $\to$.

For $\Delta = \Gamma + A - B$, show that
$\Gamma \mid\sim A\to(B\to^{-}C)$ iff $\Gamma \mid\sim B\to^{-}(A\to C)$.
<u>Left to right</u>:
Given $\Gamma \mid\sim A\to(B\to^{-}C)$.
By $R\to$, this holds iff $\Gamma, A \mid\sim B\to^{-}C$.
r) By $R\to^{-}$, this holds iff $\{\Gamma,A\}-B \mid\sim C$.
We are trying to show that $\Gamma \mid\sim B\to^{-}(A\to C)$.
By $R\to^{-}$, this holds iff $\Gamma - B \mid\sim A\to C$.
r') By $R\to$, this holds iff $\Gamma - B, A \mid\sim C$.
(r) and (r') are equivalent.
This will work to show the other direction, too.

So I think I can show path-independence for the system that includes both kinds of conditional (and the only negation we need).

Would like to show further that the iterated conditionals of the same flavor can be swapped for single conditionals (of that flavor) with conjunctive antecedents:
$\Gamma \mid\sim A\to(B\to^{-}C)$ iff $\Gamma \mid\sim B\to^{-}(A\to C)$ iff $\Gamma \mid\sim (A\&B)\to^{-}C$.

Start by showing this for upward conditional:
$\Gamma \mid\sim (A\&B)\to C$ iff (by $R\to$) $\Gamma, A\&B \mid\sim C$ iff (by L&) $\Gamma,A,B \mid\sim C$ iff (by $R\to$) $\Gamma,A \mid\sim B\to C$ iff (by $R\to$) $\Gamma \mid\sim A\to(B\to C)$. And we could go from $\Gamma,A,B \mid\sim C$ to $\Gamma,B \mid\sim A\to C$ and then to $\Gamma \mid\sim B\to(A\to C)$, as well.

For downward conditional:
$\Gamma \mid\sim (A\&B)\to^{-}C$ iff (by $R\to^{-}$) $\Gamma-(A\&B) \mid\sim C$.

But now Γ need not contain A&B. It might just have A, B in it.
And if has all three, subtracting the conjunction does not subtract the conjuncts.
This perhaps shows that the rules $R\rightarrow^-$ and $L\rightarrow^-$ rely on an ill-defined notion of "subtracting" a sentence from a premise-set. If not, then it is *very* lexically sensitive: only the conjunction/disjunction is removed, nor the conjuncts or disjuncts.
I think this is a workable understanding.
But it will *not* support substituting single downward conditionals with conjunctive antecedents for long, multiply embedded conditionals. We still will get each representation of a path from Γ to Δ as a pair of conditionals—but they will be *long* conditionals.

So the result above shows that we can*not* aggregate upward conditionals by conjoining their premises, and likewise for the downward ones, so as to end with just two conjoined conditionals, one of each flavor, and each having long conjunctions for antecedents.
We can do that for the upward conditionals, but not the downward ones.

I still need to show that the order of upward and downward conditionals does not matter when we mix them. (I've only shown path-independence for sequences of conditionals of the same "flavor".)
Want to show that i) $\Gamma \mid\sim A\rightarrow(B\rightarrow^- C)$ iff ii) $\Gamma \mid\sim B\rightarrow^-(A\rightarrow C)$:

i) $\Gamma \mid\sim A\rightarrow(B\rightarrow^- C)$ iff $\Gamma, A \mid\sim B\rightarrow^- C$, by $R\rightarrow$.
$\Gamma, A \mid\sim B\rightarrow^- C$ iff iii) $\mathbf{\Gamma\text{-}B, A \mid\sim C}$, by $R\rightarrow^-$.

ii) $\Gamma \mid\sim B\rightarrow^-(A\rightarrow C)$ iff $\Gamma\text{-}B \mid\sim A\rightarrow C$, by $R\rightarrow^-$
$\Gamma\text{-}B \mid\sim A\rightarrow C$ iff iii) $\mathbf{\Gamma\text{-}B, A \mid\sim C}$, by $R\rightarrow$.

So that is as it should be.

I conclude that the logically complex consequences of every Γ encode the consequences of every Δ.
Suppose $\Delta\text{-}(\Delta\cap\Gamma) = \{X_1 \ldots X_n\}$, the set of sentences in Δ but not Γ, and
$\Gamma\text{-}(\Delta\cap\Gamma) = \{Y_1 \ldots Y_m\}$, the set of sentences in Γ but not in Δ.
Then $\Delta \mid\sim A$ iff $\Gamma \mid\sim Y_1\rightarrow(Y_2\rightarrow(Y_3\rightarrow \ldots Y_m)\ldots)) \rightarrow (X_1\rightarrow^-(X_2\rightarrow^-(X_3\rightarrow^- \ldots X_n)\ldots)) \rightarrow^- A)$.

Problem:

As it stands, adding $R\rightarrow^-$ and $L\rightarrow^-$ to NM-MS is not conservative, *if* $R\rightarrow^-$ *is reversible*, because of an interaction with conjunction. (There might be a corresponding issue with disjunction, though I don't see it yet.)

Ulf points out that the rules I offer above permit the following derivation (I've generalized it a bit):

1. Suppose $\Gamma \mid\sim C$, for Γ containing neither A nor B.

2. Then Γ,B |~ B$\rightarrow^{-}$C, by R$\rightarrow^{-}$, since by hypothesis Γ-B = Γ and Γ|~C.
3. So Γ,A,B -A |~ B$\rightarrow^{-}$C, since A$\notin\Gamma$, so Γ,A,B -A = Γ,B, and by (2) Γ,B |~ B$\rightarrow^{-}$C.
4. So Γ,A,B |~ A$\rightarrow^{-}$(B$\rightarrow^{-}$C), by R$\rightarrow^{-}$.
5. So Γ,A&B |~ A$\rightarrow^{-}$(B$\rightarrow^{-}$C), by L&.
6. So Γ,A&B, -A |~ B$\rightarrow^{-}$C, by reversing R$\rightarrow^{-}$.
7. Γ,A&B, -A -B |~ C, by reversing R$\rightarrow^{-}$.
8. So Γ,A&B |~ C, since A$\notin\Gamma$,A&B and B$\notin\Gamma$,A&B.
9. So Γ,A,B |~C, by reversing L&.

In this way, we have derived Γ,A&B |~ C from Γ|~C, and indeed, Γ,A,B |~ C, for arbitrary A,B. That need not have held before introducing $\rightarrow^{-}$.

So *if* R$\rightarrow^{-}$ is reversible, its addition to NM-MS is not conservative (given reversibility of L&).

Some Observations:

i. We saw another manifestation of this same issue, I think, in not being able to trade chained downward conditionals for single downward conditionals with conjunctive antecedents. There the issue was that removing a conjunction A&B from a premise-set is not equivalent to removing its conjuncts from that premise-set. Here we see the effects of the converse: removing conjuncts from a premise-set does not remove the conjunction.

   Note that in multisuccedent systems, including NM-MS, what is true of conjunction on the LHS is usually true of disjunction on the RHS. For instance, both of these are multiplicative in NM-MS, to avoid forcing monotonicity, while the right-hand rule for conjunction and the left-hand rule for disjunction are additive. So we might expect an analogous difficulty for disjunctions in the consequences as we see here for conjunctions in the premises.

ii. This argument does not address the reversibility of the proposed left rule for the downward conditional. Do corresponding issues arise there?

iii. The right rule is the one that defines the expressive role of the downward conditional. If it cannot be amended so as to be made reversible, while maintaining the spirit of the original right rule, then we must either give up the monadological project, give up the idea of implementing it with a downward conditional, or seek to pursue monadologicality with a downward conditional that is not (right-) reversible.

iv. Adding a connective nonconservatively is not really an option.

v. Giving up reversibility (just of R$\rightarrow^{-}$, also of L$\rightarrow^{-}$?) might be an option.

We would not have Dan's expressive completeness result for NM-MS + $\rightarrow^{-}$, but could no doubt live with that, since NM-MS would still have it.

The question would be how much of monadologicality could be recovered without reversibility of R$\rightarrow^{-}$. It seems as though in

Target: Δ |~ A iff Γ |~ $Y_1\rightarrow(Y_2\rightarrow(Y_3\rightarrow\ldots Y_m)\ldots)) \rightarrow (X_1\rightarrow^{-}(X_2\rightarrow^{-}(X_3\rightarrow^{-}\ldots X_n)\ldots)) \rightarrow^{-}A)$.

It *might* be that we would only have to give up the 'if' direction. For without reversibility, it seems we could still go from Δ |~ A to Γ |~ $Y_1 \rightarrow (Y_2 \rightarrow (Y_3 \rightarrow \ldots Y_m)\ldots)) \rightarrow (X_1 \rightarrow^{-} (X_2 \rightarrow^{-} (X_3 \rightarrow^{-} \ldots X_n)\ldots)) \rightarrow^{-} A)$.
If so, the question is how we would have to qualify the monadologics ideology to accommodate this constraint.

Idea (shot in the dark):
The difficulty arises only when we consider subtracting from premise-sets that include logically complex sentences, paradigmatically, conjunctions.
But NM-MS by itself (before we introduce the downward conditional) supports Dan's expressive completeness result (representation theorem).
So we can represent what is happening with (and what follows from) premise-sets involving logically complex sentences entirely in terms of (constructably specifiable) constraints on sets of logical atoms (and their reason relations).
Might it be possible to avoid the problem Ulf points to by

i. Reverting to the sets of atomic sentences, and their relations, in the base vocabulary, associated with each set of logically complex sentences by Dan's result,
ii. Do what we need to do with downward conditionals at that level—where the issues about logically complex sentences do not arise (note: claim that needs to be clarified and demonstrated).
iii. Then return to draw conclusions about the consequences of premise-sets with logically complex premises.

Could we in this way evade the issue presented by conjunctions (and perhaps other logically complex sentences)?

e) **Significance of monadologics for database management: turning online updating into offline updating.**

But, at the risk of moving from the sublime to the mundane, the point with which I wish to close is that there is potentially great utility today for Leibniz's ideas of holographic and indeed monadological representation of information.

A hologram differs from an ordinary photograph in that each part of the hologram carries information about the whole image. So where if one cuts of the corner of a photograph, removing 10% of the representing, one loses 100% of the information about 10% of the represented scene, doing the same to a hologram removes 10% of the information about 100% of the scene, which just becomes a bit lower-resolution. In a monadological representational system, each representing unit carries 100% of the information carried by all of them—though each from a somewhat different perspective. We can use the paired conditional logical model systematically to create holographic, monadological *database structures*.

A database *structure* is the combination of a *universe* of possible databases together with an *inference engine*. A particular universe of possible databases is defined by a language, thought of as a set of sentences (which can be labels of record-structures as complex as one likes). The possible databases relative to that language are just all the possible sets of sentences drawn from it (its powerset). The inference engine is a function that maps each possible database onto a larger set of sentences that are the *consequences* of what is explicitly entered in the database. Those entries are the *explicit* content of the database. What can be extracted from it, added to it, by the application of the inference engine is the *implicit* content of the database: literally, what is *implied* by it. Such an inference engine allows a suitable query system to use the database to answer questions that go beyond what has been explicitly entered into it as data.

Given a universe of all the possible databases generated by a language and the consequence relation defined by an inference engine on that universe, we can use our rules for conditionals to logically extend each database. Doing that will codify the whole material consequence relation of the inference engine into the content of each extended database. A complex inference engine has been traded in for more data, which represents what follows from *all* the possible databases. So then using an extremely simple purely logical inference engine, the same for each, each database can be queried not only about what is implicit in *its* content, but also about what is implicit in the content of all of its variants that would result from adding further information to the database, or rejecting some information that had been stored there. Each database then includes the whole consequence relation, and all the information from *all* the possible databases. The information is now stored holographically, indeed monadologically.