September 28, 2022

Week 5 Notes

Plan:

1) **Recap**:

2) More on the Structure of **Implication**:

a) Implicational Persistence, upward arrow, MSFs (MRs), ranges of subjunctive robustness, regions of monotonicity, nonmonotonicity as a modality (RSR rather than necessities=monotonicities), WCM, WCT, amalgamating premise-sets as an alternative to "weighing" reasons.

b) Relation of three grades of transitivity (MC and CT, mixed and shared context Cut) to three grades of monotonicity (MO, CM). Table. Ulf's two arguments.

c) Punchline: **hysteresis of rational explicitation** as the alternative to closing under consequences.

3) More on the Structure of **Incompatibility**:

a) Incompatibility and incoherence. Analogues of implicational structure for incompatibility by using $\perp$ formulation (or empty RHS).
b) Theoremhood and empty LHS, why we don't care about theoremhood in *this* sense in nonmon setting, what sense we *do* care about theoremhood is: being a consequence of *every* premise set.
c) Incompatibility is the basis for bilateralist understandings of implication.
d) Persistent incompatibility. Recipe for turning nonmonotonic implications into nonmonotonic incompatibilities-incoherences. (Cf. arg last time about connection between failures of MO and failures of CT, in re (2a)).

4) On the Structure of **Interactions** between the two sorts of Reason Relations:
a) Explosion. EFQ and EFFQ. Relation to WCM, WCT.
b) **Implicitly *in*coherent coherent premise-sets**.
c) Implicitly coherent i*n*coherent premise-sets.

5) Conclusion.

<u>Recap</u>:

Last week:

1. Two reason relations ("broadly inferential" relations).
2. One symmetric, the other not.
3. For implication: Not MO, not CT, not CM.

4. For recap:

The pragmatic sense of "implicit commitment" is the sense in which being precluded from entitlement to reject/accept is being *implicitly* committed to accept/reject.
That is the sense in which commitment to accept the *explicit* content of a premise-set commits one to accept its (rationally) *implicit* content—that is, what that premise-set *implies*.
(A corresponding thing can be said about incompatibility.)
This connection between implicit *commitment*, on the *pragmatic* side, and implicit *content*, on the side of reason relations (in a *rational* MV, an MV for specifying reason relations) is an important artifact of the transition from practices of reason*ing* to reason relations among claim*ables* (what can be accept*ed*/reject*ed*) that was the topic of Pragmatics II in Week 3.

5. Last week I talked about the structure of *incompatibility* (symmetric) and of *implication* (CO, but not MO, or even CM, and not CT, in general or globally).

I pointed out that there is a version of MO for incompatibility.
Though I did not say so, there is also a version of CM for incompatibility.
It says that although one cannot in general be sure that an incoherent set cannot have coherent supersets (MO might not hold for incompatibility), it might be the case that any set that is *persistently* incompatible with a set Γ has special properties. It is immediately obvious that adding to an incoherent set anything that is *persistently* incompatible with it results in a set that is persistently incoherent.

I would call (and have called) reason relations "broadly inferential" relations. I don't here because I'm trying to mark the practices/relations (a version of ing/ed) distinction by infer(ring)/implication.

I'm not aware of people having addressed most of the topics I'll be talking about today: the structure of open reason relations.
I conjecture that philosophers and logicians by and large have not thought about how (for instance, but the principal case) nonmonotonic implication relations work. I think that if they got in the vicinity of these issues at all, they thought they would produce a nonmonotonic logic, and then let *it* tell them about nonmonotonic consequence relations.
But, as we'll see, there are a lot of issues that don't depend at all on the details of how you implement or specify such relations.

I. <u>More on the Structure of **Implication**</u>:

a) Implicational Persistence, upward arrow, MSFs (MRs), ranges of subjunctive robustness, regions of monotonicity, nonmonotonicity as a modality (RSR rather than necessities=monotonicities), WCM, WCT, amalgamating premise-sets as an alternative to "weighing" reasons.

I argued last time, to begin with, that we should not understand implication in general to be *monotonic*.

We should think of the reasons we ordinarily give for claims as at least in principle *defeasible*.

That is, we should acknowledge the possibility that there are some potential auxiliary hypotheses or collateral premises such that *if* we added them to the premises we actually offered as reasons, the conclusion we sought to give reasons for would no longer follow.

Maybe, we can hope, that there is no set of additional premises *we would accept* (take to be *true*) that has this property.

Maybe only outlandish scenarios—maybe even only globally *skeptical* ones—would infirm the implication.

But we should acknowledge that there *are* such possibilities: circumstances that *if* they obtained would infirm the implication, perhaps even rule out the conclusion.

And similar considerations apply to the monotonicity of *incompatibility*.

Just because some set of premises provides reasons *against* a conclusion, it does not in general follow that every superset of that set of premises similarly rules it out.

Suppose Γ|~A, in the sense that the set of premises Γ gives good reason *for* concluding A—accepting A *for that reason*.

If all you *knew* (believed, were committed to accept) was Γ, or all you had *supposed* was Γ, then the proper conclusion would be to accept A.

Now **suppose that Δ#A**, in the sense that the set of premises Γ gives good reason *against* A—reason to *reject* A *for that reason*.

If all you *knew* (believed, were committed to accept) was Δ, or all you had *supposed* was Δ, then the proper conclusion would be to reject A.

Is it *impossible* that both these should be true?

Should we limit reason talk to exclude this possibility *a priori*?

If not, then we must reject monotonicity of the two reason relations: implication and incompatibility.

For monotonicity of implication commits us to **the structural metainference**:

<u>Γ|~A</u>
Γ,Δ|~A.

And monotonicity of incompatibility commits us to **the structural metainference**:

<u>Δ#A</u>
Δ,Γ#A.

But we don't know *in advance* (a reason-relation sense of "*a priori*") which of these conclusions we ought to draw—if either.
This sort of circumstance *can* occur, and *if* it does, there is no *general* recipe for resolving the situation.

Here I want to introduce an important idea.
It is important because it incorporates a change of perspective on such questions.
Suppose we are given a set of fully specified reason relations.
That is, relative a language L, thought of just as a set of sentences (or sentence letters), suppose that we *stipulate* all the good implications and incompatibilities.
So for every possible premise-set (for the moment we can think about just the finite ones), we settle what implications with those premises hold—what sentences those premises genuinely provide reasons for, what conclusions actually follow from them—and what is incompatible with those premises.
We can represent all these stipulations in **an ordered triple of sets <L,IMP,INC>.**
For simplicity, we can restrict ourselves to the single-succedent case, where what goes on the right hand side of the implication and incompatibility signs is a single sentence.
Then IMP, the set of implications, is a set of ordered pairs of a premise-set of sentences and a conclusion sentence. IMP={<S,A>: S⊆L and A∈L}.

Now we *could* represent the set of incompatibilities in just the same form.
But we have pointed out that it is a basic structural feature of incompatibility relations, as opposed to implication relations, that incompatibility is *de jure* a *symmetric* relation.
(We speculated about the deep pragmatic reasons for this structural feature.)
We can build that symmetry of incompatibility into our representation of reason relations by taking INC to be, not a set of ordered pairs of premise-sets and conclusions, but just as a set of *sets*: the set of all the *incoherent* sets.

Incoherent sets are to be understood as such that each of their elements is incompatible with set containing all the rest.
If S∈INC and A∈S, then S-A#A.
(In the multisuccedent generalization, Γ#Δ iff (Γ∪Δ)∈INC.)
So INC={S⊆L}.
We can call such a triple <L, IMP⊆Lx 𝒫(L), INC⊆𝒫(L)>, where L is a set of sentences and 𝒫(L) is its powerset (the set of subsets of L), a **"material rational frame."**
For historical reasons, I will abbreviate this "**MSF**" for "material semantic frame."
(That's what we've always called these structures in the ROLE group. But it would be getting far ahead of my story here to introduce these as specifically *semantic* structures.
So let's just use "MSF" and not worry about where the 'S' comes from.)

Now we can talk about the structure of implication and incompatibility relations in particular MSFs. To deny that implication should be understood always to hold monotonically is to acknowledge that at least for some MSFs codifying legitimate reason relations, it is not the case that whenever $\Gamma|\sim A$ is a good implication, so is $\Gamma,\Delta|\sim A$.

Of course, **to deny that monotonicity does not hold *globally*—everywhere, in every licit MSF—is not to deny that it ever holds *locally*.**

Suppose that in some particular MSF it is the case that for some Γ and A, not only is $<\Gamma,A>$ in IMP, that is, $\Gamma|\sim A$, but also, for every $X\subseteq L$, $\Gamma,X|\sim A$ ($<\Gamma\cup X,A>\in IMP$).
$\Gamma|\sim A$ holds no matter what additional premises one adds to Γ.

Then we can say that **not only does the implication $\Gamma|\sim A$ hold, but it holds *persistently*** (that is, monotonically). There is no set of sentences in L that serve as defeaters of this implication.
We can **mark such an implication with an upward arrow: $\Gamma|\sim^{\uparrow}A$**.

If **every** implication in an MSF **holds persistently**, then the MSF is has a ***monotonic* implication** relation.
(We can define the corresponding notion for incoherence, and so for incompatibility.
If, according to an MSF, a set $\Gamma\cup\{A\}$ is incoherent, and so are all of its supersets, then it is *persistently* incoherent. Then we can write $\Gamma\#^{\uparrow}A$.)

Now it is important that even an MSF in which not *all* implications (or incompatibilities) are persistent, *some* can be.
That is, reason relations as represented by **an MSF can have local *regions* of monotonicity**.
These will be implications such as "The cloth is scarlet" |~ "The cloth is red" that can*not* be defeated or infirmed by the addition of further premises—unlike, say, "Tweety is a bird" |~ "Tweety can fly."
In general, MSFs will have a *mixture* of persistent implications and incompatibilities, and defeasible, nonmonotonic ones.

The result is that in any particular MSF, **persistence** (local regions of monotonicity) of reason relations shows up as **a kind of *modality*.**
Persistent goodness of an implication is a kind of *necessity*.
If Γ *persistently* implies A ($\Gamma|\sim^{\uparrow}A$) then we could say that it *necessarily* implies A, that is, no matter what—no matter what other premises are supposed.

Whereas if Γ implies A, but **only defeasibly, that implication is in this sense merely *contingent***.
It holds so long as certain other conditions are *not* met.
As we will see, it turns out that relative to an MSF one can introduce a modal operator, a bit of logical vocabulary, explicitly to mark reason relations that hold persistently:
$\Gamma|\sim\Box A$ iff $\Gamma|\sim^{\uparrow}A$.

But that is a story for later on, when we pursue the idea of various kinds of modality as marking regions of local structure in reason relations.
For now it is enough to remark that **what we usually care about is not just whether or not an implication (or incompatibility) is *persistent*, but what its specific *range of subjunctive robustness* is.**
That is, where an implication is not persistent, as in

The hungry lioness notices the wounded gazelle

implies

The hungry lioness will chase the wounded gazelle.

we know that *some* further collateral premises would infirm the implication.
After all,

The hungry lioness notices the wounded gazelle and she is immediately struck by lightning

does *not* imply

The hungry lioness will chase the wounded gazelle.

An essential aspect of understanding the original implication, and what it tells us about hungry lionesses and chasing wounded gazelles, is understanding the difference between further conditions under which the implication *would*, and further conditions under which it would *not* remain good.

All this information is encoded in an MSF.
For it tells us of *every* set of premises, and *every* candidate conclusion, whether those premises imply that conclusion. So it tells us *which* of the supersets of Γ continue to imply A and which do not. We can read off of the MSF all the ranges of subjunctive robustness of all the good implications (and, as will become important when we start actually to do semantics in terms of reason relations, for all the candidate implications that are *not* good, what further premises one would need to add to them to *make* them good).

From this point of view, the traditional insistence on treating only persistent implications as good ones (that is, enforcing global monotonicity of reason relations) is **like restricting oneself only to claims that are *necessarily* true, and ignoring all the claims that, while true, are only *contingently* true.** The semantically important idea of truth conditions would go missing.
Ranges of subjunctive robustness **are to reason relations what *truth conditions* are to sentences**.

There are different things to mean by 'implies' and 'incompatible-incoherent.'
The tradition chooses to restrict its attention to the sense I am picking out as *persistently* implies/incoherent.
But we can get more generality by treating that as just a special case.

The only reason not to is if the more open structures of reason relations were either unrecognizable as such, or were totally intractable conceptually and formally.
But neither of those is so.

Distinguishing, among the good implications (and incompatibilities) those that are *persistently* good and those that are only *defeasibly* good brings into view more structural conditions of reason relations that we can investigate.

I pointed out last time that in addition to full monotonicity (MO), we can consider a weaker, fallback structural condition on implication (and incompatibility), what I called "Cautious Monotonicity" (CM).
This is the condition that is met when a good implication can be defeated by the addition of *some* further premises, it cannot be defeated by the addition of any further premises that it already implies. Its own rational consequences don't defeat the implication.
In any given MSF, some implications might have this feature, and others not.
Last week I argued that, as we can now put the point, not only should we not restrict ourselves to MSFs all of whose implications are persistent—and so satisfy full monotonicity (MO)—but we should *also* not restrict ourselves to implications that satisfy CM.

For, **together with** a standard transitivity condition **CT, Cautious Monotonicity entails** that **explicitation**—explicitly acknowledging and treating as a further premise some consequence of the premises—**would be inconsequential**: it would make no difference to what follows from a premise set.

We are now in a position to **formulate a structural principle weaker than CM**.
What I will unimaginatively call "***Weak* Cautious Monotonicity" (WCM**) is the condition that although in a particular MSF an implication Γ|~A might be defeasible, and it might even be defeasible by explicitating some of its other consequences, it is *not* defeasible by explicitation of any of its *persistent* consequences.

WCM: $$\frac{\Gamma |\!\sim^{\uparrow} A \quad \Gamma |\!\sim B}{\Gamma, A |\!\sim B.}$$

There is a corresponding weaker version of the transitivity principle that is dual to CM:

WCT: $$\frac{\Gamma |\!\sim^{\uparrow} A \quad \Gamma, A |\!\sim B}{\Gamma |\!\sim B.}$$

By definition, no additional premises can turn either an implication or an incompatibility from a *persistent* one to one that is *not* persistent.
But it could do the opposite: turn a non-persistent into a persistent one.
These are what I am calling "*solidifiers*."

They are a kind of opposite of *defeators*, but at the level of *persistence* rather than just *implication* or *incompatibility*.
At the level of implication/incompatibility, what contrasts with *defeaters* (defeasors) is something that turns a bad implication into a good one, or a set that is not incompatible into one that is.
These are what Dan appeals to in his semantics.

Principles Relating Explicit and Implicit Content

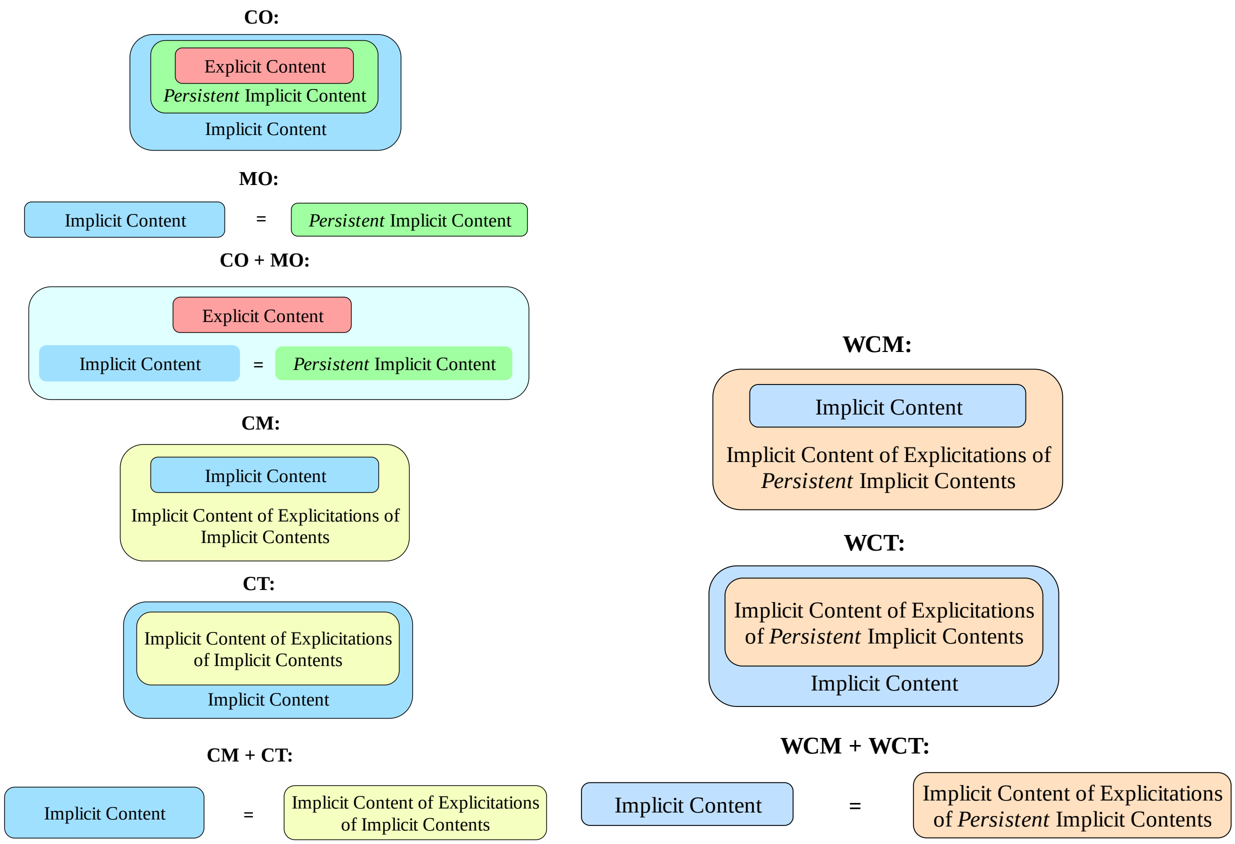

I started with the question of how we should think about a case where (as we can now say, according to some MSF) Γ provides reasons *for* A (implies it) and some other premise set Δ provides reasons *against* A (is incompatible with it).
The tradition invites us to think of this situation as calling for the ***weighing* of the *pro tanto* reasons Γ and Δ. Which provides the *stronger* or *better* reason**?
We are now in a position to think about such situations differently.
What we are asking is what follows about A if we *amalgamate* the premise sets?
Is A implied by or incompatible with Γ∪Δ?
Any particular MSF will encode an answer to that question.

Is Δ in the range of subjunctive robustness of Γ|~A?
Is Γ in the range of subjunctive robustness of Δ#A?
In each case we are asking about the *neighboring regions* of these implications and incompatibilities, in the MSF, though of as comprising all the good reason relations.
Of course, treating an MSF as an oracle in these matters doesn't make things any easier *epistemically* in assessing any particular collision of reasons.
But it turns out to make important differences in how one understands these situations.

It is if "reasons" just *means* "decisive, ***dispositive*** **reasons**."
But jurisprudence acknowledges also (and could it not?) merely ***probative*** **reasons**.
Not just reasons "all things considered," or "on balance," but reasons *pro tanto.*
Prima facie reasons, which must be assessed in concert with all the other reasons, pro and con.
Perhaps we cannot think about such reasons systematically, formally, logically, or mathematically. Perhaps.
But should we concede that from the get-go?
Only if we are constrained by what we believe our formal apparatus can handle, and let that decide, rather than devising ways to apply that apparatus to the case that really matters.
Do we look for the keys where the light is good? Or where the keys are?

To give just one example of how thinking in terms of amalgamation of premise sets instead of weighing of reasons can make a difference:
John McDowell strenuously objects to the use of "hydraulic" imagery of "weighing" or "assessing the strength or power" of different reasons.
He thinks we are misled if we think in terms of the "overridingness" of moral reasons.
That is to think of what Kant called "categorical" imperatives as just a really strong species of "instrumental" imperatives, and leads otherwise astute theorists of practical reason such as Philippa Foot to wonder about the possibility of *weak* moral reasons, that can be outweighed by strong prudential ones.
In his view, **moral considerations should be thought of as "*silencing*" other considerations, not "outweighing" them**.
Thinking in terms of local regions of persistence of reason relations in an MSF gives us a way to understand clearly what is at issue in such claims.

b) Relation of three grades of transitivity (MC and CT, mixed and shared context Cut) to three grades of monotonicity (MO, CM). Table. Ulf's two arguments.

We see that there are different grades of monotonicity, and so of nonmonotonicity.
I suggested that reason relations that fail not only MO but CM be called "hypernonmonotonic."
We just saw that even hypernonmonotonic implications can satisfy a weak kind of monotonicity, namely if they are indefeasible by the explicitation of their *persistent* consequences.

Below even this level of structure is CO, Containment.
It says that explicit content is always also implicit, that premise sets always imply their own elements.
We can see that it is a kind of monotonicity, for it says that inferences of the form Γ,A|~A are always persistent: that Γ,A|~$^{\uparrow}$A.

In fact, there is an intimate connection between these different grades of monotonicity and different transitivity principles.
(I put into the "Supplementary" readings a link to an article by Dave Ripley that offers a comprehensive survey of transitivity principles. From our point of view, it is unfortunate that he does so only within the context of a global commitment to full monotonicity. So he doesn't address the issue of the interaction of transitivity and monotonicity that I want to introduce here.)
In general clamping stronger transitivity principles onto weakly monotonic consequence relations can result in forcing stronger monotonicity structure.

For instance, a very strong classical transitivity structural principle is Mixed-Context Cut, or MC:

Mixed Cut (MC): $$\frac{\Gamma|\!\sim A \quad \Delta,A|\!\sim B}{\Gamma,\Delta|\!\sim B.}$$

It is "mixed" because the two contexts, Γ and Δ are combined in the conclusion.
The idea is that the second premise says that all Δ needs to imply B is A.
Γ gives us A (implies A).
So instead of adding A, we add something, Γ, that implies it.
Then we can "cut" out the A, and conclude that the combination of something, Γ, that implies A and something, Δ, that with A implies B must imply B.
This is not implausible.
As an example, suppose first body of evidence tells us there was a cat in the room,
and the second body of evidence says that, assuming there was a cat, a cat did the damage.
Then MC says the two bodies of evidence, taken together, tell us (imply) that a cat did the damage.

But if this holds in general, then so does monotonicity.
It assumes that Γ can't be evidence against B.

In the context of even the very weakest monotonicity principle, CO, Mixed Cut forces full monotonicity.

Monotonicity (MO): $$\frac{\Gamma|\!\sim A}{\Gamma,\Delta|\!\sim A}$$

(The conclusion could just be Γ,B|~A, since we can build up finite Δ step by step.)
For suppose we have Γ|~A and C is some element of Γ.

Then Γ,C|~A, since $\Gamma \cup \{C\} = \Gamma$.
Δ,C|~C, by CO.
By MC, we can then "cut" C from the premise-set of Γ,C|~A, and the conclusion of Δ,C|~C, and "mix" what's left over on the premise side, to get Γ,Δ,C|~A. Since by hypothesis $C \in \Gamma$ it follows that Γ,Δ|~A. But then CO and MC have taken us from Γ|~A to Γ,Δ|~A for arbitrary Δ. That is just MO.
(Thanks to Dan Kaplan for this argument.)

We can line up monotonicity principles and transitivity principles in a chart:

Levels of Rational Structure	**Monotonicity Principle**	**Transitivity Principle**
Traditional Urestricted (Classical/Intuitionist)	MO	MC
Implication Restricted	CM	CT
Persistence Restricted	WCM	WCT
Membership Restricted	CO	Contraction

On the last row, the minimal monotonicity principle (rejected by relevance logicians) is:

CO: $$\frac{A \in \Gamma}{\Gamma, B |\!\sim A}$$

The analogous transitivity principle (thinking of the duality of CM and CT) would be:

$$\frac{A \in \Gamma \quad \Gamma, A |\!\sim B}{\Gamma |\!\sim B.}$$

But that will hold so long as Γ is a set.
The feature that matters is what Gentzen called "Contraction":

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

You might think this is trivial, but it is very important to the French logician Jean-Yves Girard, the developer of linear logic, who says of Contraction:
"I will give anyone who affirms it two kicks. Not just two copies of one kick."

If Mixed Cut is too strong a structural transitivity principle to impose in a nonmonotonic setting, we can look at CT, which is a much weaker transitivity principle. Recall that

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

Here we have a context, Γ, that is *shared* by the two premises—by contrast to the previous *mixed* context. Now there is no evident presupposition of monotonicity needed to motivate the transitivity principle. The very same premise set, Γ, that implies A is the one that, together with what it implies, A, has the consequence B. The conclusion is that that premise set already implies the conclusion B, since it, that very same premise set, already implies what it needs in order to imply the conclusion.

Recall that I argued last week that CT has bad effects in the context of CM (the first grade of nonmonotonicity). Together, they imply that explicitation is inconsequential. And we should not build that into out assumptions about reason relations.

But things are worse than that with CT and monotonicity principles.
Now, unlike MC, CT by itself does not imply or presuppose MO.
But we don't need to add very much to it to get MO.
The argument of Ulf's that I am about to rehearse gets ahead of the arc of the course a little bit, since it involves (for the first time) *logical* vocabulary.
And officially that is not on our agenda until next week.
But looking a little bit ahead is no bad thing.
The basic idea of the *expressivist* view of logic that I will be developing is that it is the defining characteristic of specifically *logical* vocabulary, what distinguishes and demarcates it, is that it serves to make reason relations *explicit* in the object language.
That is, logical vocabulary adds to a prelogical language expressive resources sufficient to let one *say* in the logically extended language, what follows from what and what is incompatible with what in the original vocabulary—and, as a bonus, in its logical extension as well.
This claim is the basic thesis of *logical expressivism*.
For implications, which is our current topic, the idea is that *conditionals* codify implications in *sentences*, which can then themselves stand in reason relations of implication and incompatibility with other sentences.
To be properly understood as codifying implications, they must satisfy the Deduction-Detachment condition (what I'll sometimes call the "Dual Ramsey" condition):
Deduction-Detachment (DD): $\Gamma|\!\sim A \rightarrow B$ **iff** $\Gamma, A|\!\sim B$
A premise set implies a conditional just in case if we add the antecedent of the conditional to the premise set, we get a new premise set that implies the conclusion of that conditional.
This is what it means, I claim, for the conditional, in the context of a particular premise set Γ, to *say that* a particular reason relation holds. It *says that* in that context, A implies B.
Playing this expressive role, I claim, is what it *is* to *be* a *conditional*.

Aside: Though I'm not going to talk about it further today (that's for next week), the corresponding condition determining the essential expressive role characteristic of the logical notion of *negation* is this:
Negation: $\Gamma|\!\sim \neg A$ **iff** $\Gamma \# A$
A premise set implies a negated sentence just in case it is incompatible with that sentence.
By satisfying this condition, negations make it possible to *say that* the negated sentence is incompatible with a premise set.

We saw that the strong transitivity-of-implication principle of Mixed Cut builds in commitment to monotonicity.
We are now in a position to see that even the weak transitivity-of-implication principle of CT makes commitment to monotonicity hard to avoid.
For, as Ulf argues[1] (in his "Shopper's Guide"), if we assume CO, all we need to add to CT to derive MO is the expressive power of a conditional satisfying the Deduction-Detachment principle.
For we can argue as follows:

	Γ,A,B|~A	by CO.
	Γ,A|~B→A	by DD (right to left).
	Suppose Γ|~A.	
Then	Γ|~B→A	by CT, from Γ|~A and Γ,A|~B→A.
So	Γ,B|~A	by DD (left to right).

But that means that CO, CT, and DD imply MO, since they take us from the supposition that Γ|~A to the conclusion that Γ,B|~A, for arbitrary B. And that is just MO.

There are a couple of ways of thinking about the significance of this result.
We are assuming CO.
Denying it takes us into the area of relevance logics, which is rich and interesting, but not the region I am exploring in this course).
Then if we want to think about material consequence relations that are not always and everywhere monotonic, we have a choice:
We can drop the global requirement that implication relations be transitive, even in the very weakest sense, which is satisfying CT.
Or we can give up the idea of codifying implication relations *logically* by introducing conditionals, which are defined by playing the expressive role specified by the Deduction-Detachment condition.
For in the context of CO and CT, introducing the expressive power of the conditional, which lets us make implication relations explicit, forces global monotonicity on a vocabulary, even if that base vocabulary did not previously satisfy that structural condition.

We want to explore the possibility of introducing logical vocabulary with the expressive power to make reason relations of implication (and incompatibility) explicit—that is, conditionals (and negation) to base vocabularies, such as ordinary empirical descriptive vocabulary, where implications are not always monotonic. That requires, I will claim, DD, the Deduction-Detachment principle.

[1] He credits Arieli & Avron, "General patterns for nonmonotonic reasoning: from basic entailments to plausible relations" *Logic Journal of the IGPL*, 8(2) (2000), 119–148.

c) Punchline: **hysteresis of rational explicitation** as the alternative to **closure** under consequences.

I have been considering in what might I realize might seem to be excruciating detail **the fine structure of implication relations**, in particular stricter and more relaxed principles of monotonicity, stricter and more relaxed principles of transitivity, and some important interactions among them.
This matters for *logic* because I will argue that we want a logic that can express and help us work formally and systematically with reason relations *no matter how strict or loose* their structures are. And ordinary logics simply can't.
But that's next week's topic.
I want to point now to a *philosophical* consequence of the considerations I have been outlining.

Monotonicity and Transitivity are the most important dimensions of *closure* structures.
In their strongest form, MO and MC, (and in the context of CO), they entail that consequence relations satisfying them are topological closure operations.
Each set of sentences Γ has a single set of consequences Con(Γ).
Adding more premises to Γ always results in a new premise set whose consequences are guaranteed to include the consequences of Γ.
And the consequences of the consequences of Γ are just the consequences of Γ.
Extracting consequences from consequences doesn't add anything further.

This means that for any set of beliefs or commitments, **it makes sense to talk about *the* consequences of those beliefs or commitments**: the further beliefs or consequences that follow from or are implied by the original set.
This is **the *rational closure* of the original premises**, everything that they commit one to or give reasons for.
It is often though, for instance, that logic might play an important role in helping us compute or calculate those consequences.
Philosophers (for instance, Harman) have worried about "logical omniscience" (which, once we admit material consequence relations need have nothing in particular to do with *logic*), in the sense of believing all the consequences of our beliefs.
That *seems* like an obligation of rationality.
But it also seems hopelessly unrealistic.
Any set of commitments has an infinite number of logical consequences, and at least an indefinitely large number of material consequences.
Is it really reasonable to demand completeness of that sort as a condition of rationality?

But now we come to what seems to me a point of considerable potential philosophical significance.

If we simply deny, or even substantially **weaken the structural demands of monotonicity and transitivity** (moving down to lower rows on our chart), then rational consequence no longer has a **closure structure**.
There is no closed set of consequences for each premise set, which cannot be added to or subtracted from by continuing to extract consequences.

And that means that **the process of explicitation**, of explicitly acknowledging consequences of the claimables one accepts, and using them as premises in the drawing of further consequences, for which they provide reasons, **becomes an open-ended enterprise**.

There is no unique endpoint that it is guaranteed eventually to reach, if one just keeps at the process of explicitating rational consequences indefatigably and makes no mistakes.

- Acknowledging some consequences one's commitments give one reasons for can put one in a position where one no longer has sufficient reasons to draw other consequences that one *could* have drawn (that one would have been entitled to draw) from the original set.

That is what the **failure of CM** to hold globally means.

- And explicitating to begin with only some of the consequences of the original premise set might well provide reasons for new consequences, that the original premises did not directly provide reasons for.

That is what the **failure of CT** to hold globally means.
Explicitation is not inconsequential.
Making implicit (implied) consequences explicit as further premises can both add and subtract consequences, in the sense that the consequences of the new, enlarged premise-set might be different in both directions (might involve both additions and subtractions) from the consequences of the original premise-set.

Where implication does not have the full closure structure of strict monotonicity and transitivity, the process of extracting consequences becomes *path dependent*.
The order in which you proceed in explicitating the implicit content of your premises makes a real difference to where you can arrive.

The technical term for **path dependence is *hysteresis***.
Rational *hysteresis*, **hysteresis of rational explicitation** is the alternative to **rational *closure***.
It is what you get when you consider reason relations that are hypernonmonotonic and hypernontransitive—that is, that have local regions that fail both CM local regions that fail CT.

It is worth emphasizing that this path dependence of reasoning, in the sense of extracting consequences implied by prior commitments, is not at all a *psychological* matter.
And it is not (yet) a *logical* one, either.

For it does not turn on the presence in the language of specifically *logical* vocabulary.
It is a structural feature of reasons themselves.

This is an idea that has radical consequences for epistemology and the philosophy of mind, and also (of course) for the logic and the philosophy of logic: wherever the notion of reasons and rationality is important.

For I have mounted an argument that **the very *concept* of believing all the consequences of one's beliefs *makes no sen*se**.
For it *only* makes sense if we assume that material reason relations have fully closed structure: not only Containment, but at least Cautious Monotonicity and Cumulative Transitivity.
And I have argued that material reason relations do *not* have that structure.

Some of my larger purposes depend only on a weaker claim:
We should not build into our *logics* an expressive restriction that permits them to codify only reason relations with fully closed structure.
But I am now pointing to the significance of the stronger claim, which is what I have argued for so far—in advance of considering the relations between logic and reason relations generally.

I have mentioned that Harman and others are concerned to argue that the idea that rational beings are obliged to believe all the consequences of their beliefs is *implausible* for finite beings.
(The fact that it is *not* in that sense and for the same reasons implausible if applied to *commitments* rather than beliefs is one reason one might have for preferring normative-status talk of commitments over intentional-state talk of beliefs, in this context.
But we are now considering a claim that is *much* more radical than that claim.
It is that **it does not in general even *make sense* to talk about "all the consequences" of my commitments or beliefs**.
For that "all" presupposes a closure structure that includes monotonicity and transitivity in a strong sense.

But the idea of closing a set of beliefs under rational consequence is called on to do substantial work in epistemology, where it functions as an ideal of rationality.
(It is, for instance, a presupposition of Kant's notion of systematicity.)
To claim, as I am, that it is an ultimately unsustainable, implausible, and in many ways not even intelligible ideal—the closure structure is intelligible, of course, but the idea that it might governs ordinary reason relations is ultimately not, I think—to reject the rigorism of rational closure might seem to be a distinctive and objectionable form of semantic and therefore epistemological skepticism.

These epistemological concerns carry over to some areas of the philosophy of mind.
Philosophers of science are accustomed to think about the consequences of adopting a new theory in terms of appreciating all its consequences.
And this is to say nothing of the role of "considering all the consequences" of *philosophical* theories, views, and programs plays in *our* everyday professional lives.

To repeat: Extracting consequences from a premise set, explicitly acknowledging them, and then extracting more of the implicit content of that set is not a process that is guaranteed either to converge or to reach the same set of consequences of consequences of consequences… in the cases where it *does* converge.
Reasoning in this sense, making explicit what is rationally implicit in a set of commitments, is in principle severely *path-dependent*.
It matters what choices one makes about what to explicitate *first*, and in general the *sequence* of explicitations one effects, where you get to.
This is the ***rational hysteresis*** **of explicitation**, hysteresis of making explicit or acknowledging explicitly what is rationally implicit in (implied by) a set of commitments.

This, in a nutshell, is why *history* is an intrinsic feature of *reason*.
That is, it is why *discursive* creatures, creatures who can *talk*, which is, I have claimed, engaging in practices of making claims and challenging and defending them, have *histories*, and not just *natures*.
Enlightenment conceptions of rationality were completely innocent of this idea.
And Tarski and Gentzen's formalization of closure structures largely insulated and excused analytic philosophers of the last 90 years or so from concerning themselves with the essentially *historical* character of reasoning.
But on the line of thought I am developing here, the historical path-dependence of consequences and incompatibilities is a deep *structural* feature of reasons as such.

d) <u>More on the Structure of **Incompatibility**</u>:

Here I should talk, as I did not last week, about **irreducibly incoherent *triads* (and so on for larger sets), and tell the story about the Sellars challenge w/res to *taste*.**
Ripe, red, blackberry.

Note that **I did *not* talk last week about the relation between <u>incompatibility</u> and <u>incoherence</u> of *sets* of sentences**.
This connection is secured by arguing for the necessary *symmetry* of incompatibility, and helps explain it.
The underlying phenomenon is that *sets* of commitments to accept are such that one cannot be entitled to commitments to accept *all* of them.
This is what I have talked about under the heading of the "holism" of entitlement, by contrast to the "atomism" of commitment.
Then it follows that commitment to accept any subset is *incompatible* with commitment to accept any other set whose union with the first is incoherent.

In single-succedent sequent calculi, we encode *incompatibility* relations in *consequence* relations by detouring through *incoherence*.
We write Γ|~⊥ ('⊥', read "perp", marks the False, or the absurd) to say that Γ is incoherent.
If Γ,A|~⊥, then Γ is *incompatible* with A: Γ#A, as I am writing it.
This builds in the symmetry of incompatibility.

We *can*, then, read off the structural features of incompatibility from those of consequence.
If consequence is monotonic, so that the structural metainference from Γ|~A to Γ,B|~A is good, then treating the perp on the left just like a sentence letter, we get that the structural metainference from Γ|~⊥ to Γ,B|~⊥ is good.
And that means that incompatibility is monotonic: if Γ is incoherent, so is any superset of Γ.
Using this notation means we can't treat the two reason relations as satisfying different structural principles: for instance, implication satisfying CM but not MO, but incompatibility being fully monotonic—or *vice versa*.
We will avoid these commitments of notation by keeping separate books on |~ and #.
Note that in multisuccedent sequent calculi, the same effect is achieved by using *empty* right-hand sides: Γ|~ .
And empty *left-hand* side is used to mark what follows from *no* assumptions (the empty set of premise). If |~A, then A is a theorem.
It is worth pointing out that although our MSFs allow good implications from empty premise-sets, in a nonmonotonic setting *this* notion of a theorem is of little interest.
For the notion of theorem that *matters* is what follows from *every* premise-set—"no matter what." In a monotonic setting, that condition is entailed by following from nothing.

But in a nonmonotonic setting it is different. What follows from the empty set might not follow from *any* populated set of premises, never mind *all* of them.
In a particular MSF, there might be some claims that follow from *every* premise set.
They are its *material* theorems.

At any rate:
We should consider nonmonotonic incompatibility relations just as we consider nonmonotonic consequence relations.

Is there a **recipe for turning nonmonotonic implications into nonmonotonic incompatibilities**?
There should be.
If Tweety is a bird |~ Tweety can fly, then in the same sense it should be that
Tweety is a bird # Tweety can't fly. There is some at least *incongruity* here.
But Tweety is a bird and Tweety is a penguin √# Tweety can't fly.
Indeed, in *this* case (but see below):
Tweety is a bird and Tweety is a penguin |~ Tweety can't fly.

Pirmin is drinking a beer # Pirmin does not drink alcohol.
Pirmin drinks beer, and the beer Pirmin drinks is O'Doul's √# Pirmin does not drink alcohol.
(O'Doul's is a nonalcoholic brand of beer.)
And in this case
Pirmin drinks beer, and the beer Pirmin drinks is O'Doul's √|~ Pirmin does not drink alcohol.
(Maybe he drinks other kinds of alcohol.)

What is going on in these examples is that if Γ|~B and B#C, then often (usually? always?) Γ#C.
But if Γ,A√|~B (A defeats the implication of B by Γ), then it often (usually? always?) also defeats the incompatibility of Γ with C.
(Q: When would it not? ◊A: When Γ#A.)

Notice that in both the Tweety case and the Pirmin case, the additional premise that defeats the incoherence is *compatible* with the premise set.
In the first case it implies the conclusion of the implication all on its own, in the second it does not imply the conclusion (alone or with the rest of the premises).

What about self-incompatibles?
We assume that the whole language L is incoherent, so that there are some incompatibilities.
Should we exclude self-incompatible sentences?
All incoherent sets are self-incompatible, so why deny the intelligibility of that notion for sentences?
Perhaps this is what paradoxical sentences are.

And notice that once we have allowed them, the failure of monotonicity of incompatibility means that it by no means follows that every premise-set that *contains* such a self-incompatible sentence (or set of sentences), whether *explicitly* or *implicitly*, is itself incoherent.
That is just the consequence that nonmonotonicity (or hypernonmonotonicity) blocks.
So: don't give up the search for auxiliary hypotheses that, when added to liar sentences or Curry sentences (which are liar sentences for inferentialists), yield supersets that are coherent.

e) <u>On the Structure of **Interactions** between the two sorts of Reason Relations</u>:

This is about the fundamental structure of *reasons*.
In particular, about **the relation between rational *inclusion* (implication) and rational *exclusion*** (incompatibility).
In some sense, we have been worrying about this since Aristotle.
But some basic lines of thought have not been sufficiently pursued at this highest level of structural abstraction.
(Hegel did what he could. These are his "mediation" and "(determinate) negation"—the inhalation and exhalation of inclusion (advancing by drawing further conclusions) and exclusion (securing by critically ruling out errors).
This is what shows up for us, in a pragmatics-first order of explanation, as *defending*, by giving reasons *for,* and *challenging*, by giving reasons *against*.

Pursuing our pragmatics-first order of explanation, we have seen **a fundamental connection between implication and incompatibility-incoherence**.
For our understanding of these two reason relations, we deepened and developed Restall and Ripley's *bilateral* understanding of what is expressed by sequent turnstiles.
That sort of bilateralism has two basic ideas:

i. The left-hand, premise side of the turnstile concerns practical attitudes of *acceptance* of claimables, expressed by speech acts of *assertion* of them,
while the right-hand, conclusion side of the turnstile concerns practical attitudes of *rejection* of claimables, expressed by speech acts of *denial* of them.

ii. The normative significance of reason relations, paradigmatically implication, is to be understood in terms of the *incompatibility* or *incoherence* of the *position* (status, constellation of commitments) that is the *combination* of one's attitudes towards the premises and one's attitudes towards the conclusion.
Accepting all the premises is *incompatible with, rules out* denying all the conclusions.
Such a position is *out of bounds* or *incoherent*.

We saw that it is easy to extend **this understanding of *implication* in terms of the *incoherence*** of some acceptances with some rejections to an understanding of *incompatibility* in terms of the incoherence of a whole set of acceptances: of the union of the premises and the conclusion.
So bilateralism is a broadly *incoherence*-first approach to defining reason relations in normative pragmatic terms.

<u>**Explosion**</u>:

a) Explosion, EFQ is a general structural condition on how *incoherence* of premise set affects *implications*.
b) Once we deny MO as a *global* structural principle for incompatibility, we should distinguish *merely*, but *curably* incoherent sets, some of whose supersets are *not* incoherent, with *persistently* incoherent premise-sets.

Example: wave behavior and particle behavior are incompatible in classical mechanics. In quantum mechanics, with lots of other auxiliary hypotheses added, they become compatible.

c) That opens up the possibility that merely and persistently incoherent premse sets behave differently in implication relations.
d) EFQ and EFFQ. Might retain an explosion principle, but restrict it to *persistently* incoherent sets. Such a position would be a *via media*, splitting the difference between traditional (classical and intuitionist) endorsement of explosion as a global principle and rejecting it entirely (as relevantists do). It could claim to retain whatever was right about the traditional view, while still allowing that we can reason with at least some incoherent premise sets—that we need not simply throw up our hands and surrender. For intelligible MSFs might still involve distinguishing, for at least some incoherent sets, between things that *do* follow from them (are implied) and things that do *not* follow from them (are not implied by them. That is the minimum we need to be able use them in reasoning.

For a persistent incompatibility, suppose I believe the axioms and rules of Peano arithmetic, but never having heard of Gödel, also believe that there is a complete finite axiomatization of arithmetic. Then my commitments are implicitly incoherent—indeed, *persistently* implicitly incoherent. Does that mean that I am precluded from entitlement to deny (so, implicitly committed to accept) that I can put my unprotected hand in a flame without pain or injury? Surely no-one really thinks that.

Fallback: from EFQ to EFFQ.

I can make my implicit commitments explicit and explicitly accept that *p* and *not-p* for some claimable: say, that Peano arithmetic is complete and Peano arithmetic is not complete. This is an instance of a *persistent explicit* incompatibility, indeed *inconsistency*.
Whatever plausibility there is to EFQ (which is not much) is retained by this instance of EFFQ. At least the contagion is confined.

[BB: The last part of these notes themselves dissolve, if not into incoherence, into lack of a clear narrative line. Sorry about that.]

Looking at the *interaction* of the two sorts of reason relations:

Incompatibility-incoherence is nonmonotonic (fails the analogue of global MO) iff it can be that:

Γ#A and *not* Γ,B#A.

Incoherence is defeasible; it is not always persistent; it can sometimes be “cured” by the addition of further premises.

a) Incompatibility-incoherence fails the analogue of global CM if it can happen that Γ#A and Γ|~B and *not* Γ,B#A.

Then $\Gamma \cup \{A\}$ is *explicitly* incoherent, but *implicitly* coherent.
For it can be turned into a coherent set by explicitation.

b) Dually, Incompatibility-incoherence that fails the analogue of global CM allows that *Not* Γ#A and Γ|~B and Γ,B#A.

Then $\Gamma \cup \{A\}$ is *explicitly* coherent, but *implicitly in*coherent.
That reasons and reason relations of consequence and incompatibility can intelligibly be understood to allow these possibilities is richly philosophically suggestive.
Assuming that reason relations have globally closed structures (are monotonic and transitive), closes off the possibility of thinking consecutively about these possibilities.

I. <u>Coherent sets implying incompatible conclusions—that is, *implicitly in*coherent sets</u>:

Example:
Wave behavior and particle behavior are incompatible in classical mechanics.
In quantum mechanics, with lots of other auxiliary hypotheses added, they become compatible.

a) Persistently incoherent premise sets are incompatible with *all* their consequences—even their CO consequences.
b) Can any *curably* incoherent premise set be cured by explicitating one of its consequences? The claim that, although the incoherence of some Γ can be cured by *some* additions of new premises, it can*not* be cured by the addition of any new premises that are already consequences of Γ is a sort of analogue of CM, but for incoherence rather than implication.
c) The analogue of *hyper*nonmonotonicity for incompatibility/incoherence, then, is to allow incoherent premise-sets, some of whose consequences can, when explicitated, cure the incoherence.
d) The dual of that notion is *coherent* premise-sets some of whose consequences are incompatible with them. For explicitating those consequences flips the premise set from coherent to incoherent.
e) Together, the two principles would say that explicitation never *changes* the coherence status of any premise set. This is the analogue for incompatibility-incoherence of the *inconsequentiality of explicitation* for implication.
f) Is there an analogue for CM? Could explicitation turn a *bad* implication into a *good* one?

I could use √ to indicate that a reason relation does *not* hold, so that "Γ √|~ A" means that it is *not* the case that Γ|~A, and "Γ √# A" means that it is *not* the case that Γ#A.

I will use ⊥ to indicate incoherence, so Γ|~⊥ means Γ is incoherent and Γ√|~⊥ means Γ is coherent.

Then one possible consequence of explicitation is to turn a good implication bad:

1. Γ|~A and Γ|~B, and Γ,A√|~B.
   (Would there be any reason to require that this relation be symmetric?)

Another is that explicitation turns a bad implication good:

2. Γ|~A and Γ√|~B, and Γ,A|~B. (No issue of symmetry here.)

Analogously, explicitation might cure an incoherence:

3. Γ|~A and Γ#B, and Γ,A√#B.

This is *implicit coherence* or *compatibility* of *explicitly* incoherent/incompatible sets.

And explicitation might turn a coherent set incoherent:

4. Γ|~A and Γ√#B, and Γ,A#B.

This is Γ being **implicitly incoherent**.

That is an important status a premise set can have: if you make explicit its consequences, you find out that it is incoherent *in this sense*: *implicitly* incoherent.

Many, perhaps most, philosophy books and articles turn out to be like this.

(At least it often alleged by reviewers and critics that it is.)

We should distinguish this sort of incoherence from *explicit* incoherence—which still falls short of *logical inconsistency*—though it *implies* inconsistency (since contraries imply contradictories), and so is *implicitly inconsistent*.

Should be careful with this *incompatible, so implicitly inconsistent* argument. For *inconsistencies* are *persistently* incoherent, and incompatibilities-incoherence need not be.

If Γ#A then Γ|~¬A. If Γ|~A, then Γ is *implicitly* inconsistent.

But if one *explicitates* the implications, the result is an *explicitly persistently* incoherent set, which Γ *need not have been*. In that case, explicitation changes *curable* incoherence into *persistent* incoherence. That is a substantial change. It is the equivalent, for incompatibility-incoherence, of the *consequentiality of explicitation*.

Q: What should I call this phenomenon, of explicitation changing coherence value (by analogy to denying the "inconsequentiality of explicitation"?

And what if explicitation merely changes the *persistence* value of an implication or incompatibility?

i) **Defeasibly (or curably) incoherent premise sets**.

Γ|~⊥ (or Γ|~ or—equivalently, but less misleadingly Γ# , if we use incompatibility with the empty set to mark the incoherence of Γ), and for some A∈L, not(Γ,A#) or (not Γ#A) or Γ√#A, or Γ,A√# , or Γ,A√|~⊥.

Denying MO for |~ is allowing that it can be that Γ|~B and Γ,A√|~B.

This can be thought of as a special case of that, where B=⊥.

This is easier Tarski-wise: InCoh(Γ) and Coh(Γ,A).

Do we have any plausible examples of this?
Yes.
Once we deny MO as a *global* structural principle for incompatibility, we should distinguish *merely*, but *curably* incoherent sets, some of whose supersets are *not* incoherent, with *persistently* incoherent premise-sets.
Example: wave behavior and particle behavior are incompatible in classical mechanics. In quantum mechanics, with lots of other auxiliary hypotheses added, they become compatible.
Tweety is a bird # Tweety can't fly.
Tweety is a bird, Tweety is a penguin √# Tweety can't fly.

Another possibility, not a matter of explicitation, is this:

5. Γ|~A and Γ#A. If Γ is already incoherent, this might not be surprising: incoherent premise sets can have incompatible consequences.

But this might even happen if Γ√|~⊥, that is, if Γ is coherent:

6. Γ√|~⊥ and Γ|~A, and Γ#A.

Can I think of examples of all of these?

The method of counterexamples to implications suggests that a way to show that Γ does not imply A (that Γ√|~A) is to find a B that is incompatible with the conclusion, but not the premises.
That is a B such that B#A and Γ√#B.
For then, the thought is, Γ,B √|~ A.
But if Γ |~ A, then it should follow that Γ,B |~ A.
So if Γ,B√|~A, it follows (by contraposition) that Γ√|~A.
But even it is true that (B#A and Γ√#B) entails that Γ,B √|~ A (more on this assumption below), in a nonmonotonic setting, it does not follow that if Γ|~A then Γ,B|~A.
So we can*not* argue contrapositively that if Γ,B|~A does *not* hold, then Γ|~A also does not hold.
So the "method of counterexamples," in the form that a reason against the conclusion that is not a reason against the premises (something incompatible with A but not with Γ) infirms or defeats an implication, does not work in a nonmonotonic setting.
Further, not only does it not follow from (B#A and Γ√#B) that Γ√|~ A, it does not even follow that Γ,B √|~ A.
For it does not follow from (B#A and Γ√#B) that Γ,B#A.
This entailment would hold *only* if incompatibility were monotonic: MO#.
Maybe Γ cures the incompatibility of B with A, so that B#A but Γ,B√#A.
But what we are envisaging is more than just that incompatibility is not globally nonmonotonic.
The current argument is that *if* (and *because*) incompatibility is nonmonotonic, *coherent* premise-sets can have consequences that are *incompatible* with them.

Coherence is not preserved by implication.
Put another way, what we can now describe as
***explicitly* coherent premise-sets can be *implicitly in*coherent.**
That is just to say that their *implicit* content can include consequences that are incompatible with that premise set. So, when that *implicit* content is *explicitated*, when *made explicit* by being added as *explicit* premises, the result is an (explicitly) *in*coherent set.
That there can be *implicitly* incoherent sets of premises that are not *explicitly* incoherent is a consequence of the nonmonotonicity of incoherence-incompatibility.

The situation we are trying to think (coherently!) about is that I might write a book that is materially coherent, in that there is no set of claims explicitly made in the book that is incompatible with any other set of claims explicitly made in the book, and that the book's being coherent in that sense does *not* guarantee that there is not some *consequence* of the claims made in the book, part of its *implicit* content, which is incompatible with what *is* explicitly said in the book—and so which, if *made explicit*, would render it *in*coherent.
I might actually have written such books.
(If so, it is the job of readers and reviewers to find that out. So you tell me.)

For ordinary (nonpersistently) incoherent sets, we can still distinguish between what follows from them and what does not—and even, what they are incompatible with or not.
Now you are in trouble if your explicit commitments are incoherent.
And not in as much trouble if the incompatibility is just a matter of what is *implied* by your explicit commitments.
Of course, if you then recognize the consequences and acknowledge them, you go from the second position back to the first.
Incoherent sets imply the negations of all their explicit members.
But that does not mean that one explicitly *draws* those logical consequences (and so explicitates the negations).
If and when one does, then one is another, worse position.
Because A and ¬A are *persistently* incompatible.
Now you really might not be able to distinguish what follows from what does not.
So we have at least 3 grades of incoherence: implicit, explicit, and persistent, with explicitation of implicit incompatibilities using negation leading to *persistently* incoherent premise sets, and by EFFQ (which our logic will endorse), one loses the ability to reason in this region of the MSF, because one loses the distinction between what follows and what does not, what is incompatible and what is not.

Conclusion:

One might conclude from all these considerations, from making *explicit* what is *implied by* giving up the idea of reason relations having the full topological closure structure of monotonicity and transitivity, that the tradition had the right idea in imposing those global structural restrictions.
For the price of giving it up, of imposing *only* CO on implication and only *symmetry* on incompatibility—or even in addition the weak structure of WCM and WCT on implication and an analogue of WCM on incompatibility—might seem to be chaos: anything goes!
It might seem that the result is a position that is not only to the left of rational closure, but practically only slightly to the right of "Whoopee!"
Wouldn't we be giving up all rational intelligibility of the structureless, amorphous, gooey mess, which accordingly is unrecognizable as *reason* relations at all?
Not so and far otherwise!
It will turn out (Ulf and Dan have shown), not only that such hypernonmonotonic, hypernontransitive reason relations are intellectually and even formally tractable, but that traditional logical and semantic metavocabularies can be adapted to manage them with surprisingly small adjustments.
What looks like the Wild West can in fact be tamed, and with formal tools readily constructible from those we have inherited.
The logic NM-MS and Dan's implication-space semantics work smoothly to codify the reason relations and characterize the conceptual roles conferred on sentences by standing in those reason relations for arbitrary MSFs satisfying only CO—including those where WCM or WCT or both hold but none of the stronger MO, CM, or CT do.