

July 6, 2021

On the Structure of Reasons: Pragmatics, Semantics, and Logic

In this paper I deploy a number of arguments and report a number of results due to Ryan Simonelli (simonelli@uchicago.edu), Ulf Hlobil (ulf.hlobil@concordia.ca), and Dan Kaplan (dan.kaplan@pitt.edu), who are members of our logic working group “Research on Logical Expressivism.” I mark their contributions as best I can along the way, to indicate what they are responsible for. They should not be presumed to endorse the use I have made of their work here.

Part One: Pragmatics

I)	Semantic and Pragmatic Construals of a Fundamental Bipolarity	1
II)	Reasons and Reason Relations: Symmetries and Asymmetries	5
III)	Bilateral Pragmatic Renderings of Reason Relations	14

Part Two: Semantics and Logic

IV)	Two Routes to Semantics	23
V)	The Structure of Material Reason Relations	36
VI)	Logic and the Explicit Expression of Substructural Reason Relations	46

Part Three: The Dialogic Pragmatic Structure of Reasoning Practices

VII)	“And That’s All, Folks”: Reason Relations and the Dialogic Pragmatics of Reasoning	57
VIII)	Conclusion	69

Part One: Pragmatics

I. Semantic and Pragmatic Construals of a Fundamental Bipolarity

The Fregean semantic tradition treats the distinction between *truth* and *falsity* as conceptually fundamental. This is not to say that these notions need be taken to be conceptually primitive, in the sense of not being further analyzable—though that was Frege’s own view. But it is characteristic of this tradition to understand specifically *semantic* significance in terms of the difference something makes, its relevance to, possession of one or the other of these two “truth-values.” The basic form of semantic meaningfulness is truth evaluability.

The two truth-values are not merely different (as **square** and **red** are), but exclusively different (as **square** and **circular** are). In the central, defining cases, possessing the one truth-value rules out possessing the other. It is impossible for anything to be *both* true *and* false. Further, the two truth-values are thought of as *opposites*. Appealing to an older tradition, we could say that “true” and “false” express Aristotelian contradictories rather than merely contraries. However, at least for the Fregeans, the opposition between truth and falsity is conceptually more basic than, and presupposed by, the distinction between contradictories and contraries. Much, much later in the order of explanation, sophisticated semantic theories might find reasons to relax these structural constraints for special cases, and countenance truth-value gluts and gaps: truth-evaluable items of which one says that they are either both true and false or neither. But the intelligibility of such late-coming subtleties is understood to be rooted in and parasitic on the more basic cases where what is truth evaluable takes one and only one of the two values. The home language-game of “truth-value” talk is built on and builds in the distinction and the modally robust opposition between truth and falsity.

I have been pointing to a familiar and influential philosophical tradition that is demarcated by having *semantic bivalence*, in a broad sense, as one of its guiding ideas. It is characterized by its commitment to an order of explanation that seeks to understand *meaning* ultimately in terms of truth-values. The purest, paradigmatic application, and a strong confirmation of the power of this approach was provided early on by the clear, elegant, arbitrarily recursively extensible, bivalent truth-functional semantics it turned out to make possible for classical (Boolean) sentential *logical* vocabulary.

This bivalent semantic structure has an analogue on the side of pragmatics: that is, in theories of the *use* of linguistic locutions, rather than of the *meanings* or *contents* those locutions express. Here the starting point is *acts*, things speakers and thinkers can *do*, and the practical attitudes those acts express—in a different, pragmatic rather than semantic, sense of “express”.

The fundamental acts are *affirmation* and *denial*, expressing practical attitudes of *acceptance* and *rejection* (whether overtly and publicly, or covertly, *in foro interno*).

The acts expressing attitudes of acceptance and rejection can be as simple as saying “yes” or “no” in response to a claim or a question. These forms of affirmation and denial underscore that an essential element of the attitudes of acceptance and rejection is their opposite valence, their exclusion of one another. In the central, defining cases, one may say either “yes” or “no,” but not both. The intelligibility of the cases where we do want to affirm and deny, for instance by saying “yes and no” is parasitic on the more basic cases where the opposition holds. Affirmations and denials count as *communicating information* only when the attitudes they express are treated as ruling one another out.

The sense of exclusion that applies to practical acts and attitudes is normative. One *ought* not to perform speech acts that would express acceptance and rejection of the very same claimable—even if one *could* do so. The deontic flavor of the modality of exclusion relating practical attitudes of acceptance and rejection on the side of pragmatics contrasts with the alethic flavor of the modality of the exclusion relating truth and falsity on the side of semantics. That exclusion is understood in terms of the *impossibility* (in the central, defining cases) of the very same truth-evaluable item being both true and false.

The semantic and pragmatic oppositions are linked by the possibility of understanding the practical attitude of accepting as *taking* (to be) *true*, and the practical attitude of rejecting as *taking* (to be) *false*. Such an identification is made possible by the fact that tokenings of declarative sentences can be understood as expressing both what is semantically truth-evaluable and what one can pragmatically accept or reject (affirm or deny, agree or disagree with, assent to or dissent from, say “Yea” or “Nay” to). Indeed, declarative sentences can be functionally specified just as what can play both these roles. So construed, they are the syntactic glue that binds the semantic and pragmatic aspects of discursive practice. Together, these three form what we might call the “apophantic triad” of declarative sentences (syntactic characterization) expressing what is both evaluable as true/false (semantic characterization) and toward which one can adopt practical attitudes of acceptance/rejection (pragmatic characterization).

That those attitudes can be thought of as truth assessments or evaluations—that acceptance can be understood as taking true and rejection as taking false what is expressed by declarative sentences—opens up the possibility in principle of two orders of explanation. The semantics-to-pragmatics order of explanation would begin with a prior grasp of the concepts of truth and falsity and elaborate from that an account of what one needed to *do* to count thereby as practically *taking* or *treating* what is expressed by a declarative sentence *as* true or false, thereby accepting or rejecting it. (A potential candidate on the side of acceptance might include: using it as a premise for inferences, including practical ones issuing in intentional doings.) The

pragmatics-to-semantics order of explanation would begin with a prior grasp of the concepts of acceptance and rejection and elaborate from that an account of what property one is thereby practically taking or treating what is expressed by a declarative sentence *as*. (A potential candidate on the side of truth might include: whatever is preserved by good inferences.) If neither of these seems feasible (or desirable), one might treat the semantic and pragmatic dimensions of discursive practice as explanatorily coeval, and explain why they line up structurally as they do: true is to false semantically as acceptance is to rejection pragmatically.

How one understands these relations is a matter of some philosophical significance. For the opposition between truth and falsehood can be understood as a matter of how things are in the objective world one is talking or thinking *about*: how it is with the *represented*s one is representing in thought and talk. Whereas the opposition between acceptance and rejection applies to how one represents things to be, that is, to the subjective side of the activity of *representing*. Insofar as we are interested in how these are related to one another, we should be interested in the structural *apophantic bipolarity* that shows up in the two interrelated species of the *semantic bivalence* of true and false and the *pragmatic dimorphism* of acceptance and rejection. Both, we have observed, are related by a modally robust kind of incompatibility. The semantic distinction of truth-values is articulated by the alethic *impossibility* of the same truth-evaluable being both true and false (in the same situation and at the same time). The pragmatic distinction of practical doxastic attitudes is articulated by the deontic *impropriety* of the same claimable being both accepted and rejected (by the same subject at the same time). In either case, the *opposition* semantically between true/false and pragmatically between acceptance/rejection is crucial to one's understanding of thought and talk. The *division* into ways the world could be or opposing attitudes one could take to it is fundamental to either approach. Even if we want to allow qualifications around the edges, that in central cases the values or attitudes *exclude* one another is an essential element.

A further point on the semantic side is this. The two basic truth-values (and their associated doxastic attitudes) are complementary and co-ordinate—but not of equal status. Truth has a certain axiological priority. It is the good value, the one we want. (“Truth is beauty, beauty truth,” Keats said. Falsehood: not so much.) The goals of inquiry no doubt include both accepting truths and rejecting falsehoods, but here, too, the precedence of the positive virtue over the negative one is palpable. This is so even though rejecting falsehood is an activity that goes far beyond merely avoiding error, an end that can be achieved by mere inactivity.

There is a corresponding point on the pragmatic side. The default significance of uttering a declarative sentence is to perform an act of assertion. That is the most basic kind of affirmation, expressing the practical attitude of acceptance. Is that happenstance, or a matter merely of convenience? Is it intelligible that the most basic discursive act and attitude should

instead be denial, expressing rejection? In an influential discussion, Michael Dummett addresses this question:

Why should we not instead have had a convention whereby we were taken as uttering sentences with the intention of uttering *false* ones, not as well as the convention of assertion we do have, but in place of it? This possibility has only to be stated to be recognized as spurious.¹

He considers practices that admit in principle either reading, and argues that we could never have positive reason to adopt the denial-rejection interpretation rather than the affirmation-acceptance interpretation. That is a somewhat disappointing defense of the very strong claim he has made. It seems insufficiently appreciative of the priority of the positive pole in the “two-way” rational activities that Aristotle discusses in *Metaphysics Theta*, which Irad Kimhi placed at the center of his recent deep meditations on the topic.² It is not to the point for me to pursue the matter here, but one way we might bring the pragmatic asymmetry of attitudes into view is to consider practical reasoning. Can we really make sense of practices in which the conclusions of the most basic exercises of practical reason in deliberation consist exclusively of *rejections* of courses of conduct rather than endorsements or acceptances of them? Perhaps third-personal exercises of practical reasoning in the context of assessments of the conduct of others could be purely censorious like this. But in first-personal cases it seems plausible that more positive guidance by reasons must be available, addressing what one should (has reason to) do rather than just what one should not do.

I have not mentioned *negation*—and this might seem a culpable omission. After all, isn’t the falsity of a claimable-believable-thinkable (what is expressed by a declarative sentence) just the truth of its negation? Isn’t rejection or denial of it just acceptance or affirmation of its negation? Indeed. And large swathes of the philosophical tradition have succumbed to the temptation to bring negation into the story at the ground-level. But I think we will be rewarded by exploring a direction of explanation that understands negation in terms of the relations between the semantic distinction between truth and falsity and the pragmatic distinction between acceptance and rejection, rather than the other way around. That is the essence of the *logical expressivist* approach to demarcating logical vocabulary and concepts, which I will be motivating in what follows. I will proceed on the assumption that almost all the important features of the structure of reasons and reasoning can be made visible in advance of introducing logical vocabulary and concepts.

¹ *Frege’s Philosophy of Language* [Harper and Row, 1973] pp. 317-318, from Chapter Ten: “Asserting.”

² *Thinking and Being* [Harvard University Press, 2018], pp. 19-23, 60-61, 108-111.

II. Reasons and Reason Relations: Symmetries and Asymmetries

Understanding the pragmatic attitudes (acceptance/rejection) and the semantic values (true/false) as related by the principle that acceptance is taking-true and rejection is taking-false is made attractive by the fact that it seems that *what is* accepted or rejected and *what is* true or false are the *same* kinds of things. What one can *take* true is what can *be* true. Very austere theories might have only *sentences* as what take truth-values and are accepted or rejected. (Quine is an example of such principled austerity.) When I introduced the point above, I found it natural to talk about what is true or false and what can be accepted or rejected rather in terms of what is *expressed* by declarative sentences. The thought is that the sentences themselves are candidates for semantic evaluation and pragmatic endorsement only at one remove, in virtue of what they express. The sentences are understood as expressive vehicles for what is in the primary sense truth-evaluable and the object of pragmatic attitudes (affirmable or deniable).

Why introduce this additional complicating factor? Because different sentences get truth-valued together, and accepted or rejected together, across a variety of situations. I say “I am confused on this point,” and you say “You are confused on this point,” or “Brandom is confused on this point.” At noon I say “It is now noon,” and 10 minutes later I say “It was noon then (when I last said something).” Ludwig wrote “Wovon man nicht sprechen kann, darüber muss man schweigen,” and I say “About that of which one cannot speak, one must remain silent.” When properly understood, the sentences (tokenings in the first two examples and types in the third) in these pairs stand or fall together, semantically and pragmatically. In this sense, sentences stand in *samesaying* relations to one another. This observation is the origin of the idea that what is evaluated semantically and pragmatically is what is *shared* by the sentences that belong to the same *samesaying* equivalence class. That is what invites us to think of those sentences as all *expressing* a common *content*. That expressive relation to a common content is understood as what determines that declarative sentence tokenings, or even types, *are* *samesayings*—that they stand in the *samesaying* relation to one another.

There is another sort of covariation of truth-evaluations implicit in grouping sentences into *samesaying* equivalence classes accordingly as they get the same semantic truth-evaluation in the same situations. For we can look at grouping those situations or circumstances into equivalence classes accordingly as they yield the *same* truth-evaluation for all the sentences of some *samesaying* class. This is the origin of the idea of truth conditions (and dually, of falsity conditions), or of truth *makers* (dually, falsifiers) associated with the contents expressed by the sentences of that *samesaying* class. This assimilation is not in the first instance of sentences, but of worldly states or situations in virtue of which (*samesaying* equivalence classes of) sentences are true or false. The work of the vague Scholastic metaphor “in virtue of which” in this claim is

sometimes done by talk of what *makes* (samesaying equivalence classes of) sentences true or false. The more concrete metaphor immediately requires a qualifying commentary to the effect that the sense of “making” appealed to is not in general the *causal* sense of a temporal process that yields a result or effects a change. Moving in the other direction, the metaphor can be emptied out to the abstract form of what “determines” truth-values that invokes nothing more than a function in the mathematical sense, whose values (for some worldly arguments) are truth-values. In between is assertion of a modally robust covariation, distinguished by its support of subjunctive conditionals of the form “If the worldly situation *were* (or had been) thus-and-so, the (samesaying equivalence class of) sentences *would be* (or would have been) true (respectively: false).”

We might use the term “truth conditions” generically for the equivalence classes of worldly items playing this semantic explanatory role with respect to the samesaying equivalence classes of declarative sentences (tokenings and types). Some conceptions of truth conditions appeal to how it is with the worldly items the sentences (or the contents they express) are understood to be *about* (represent, denote, or refer to). On other conceptions, what makes sentences true are thought of as *facts*. (This is a non-Fregean sense of “fact,” since in his usage a fact just is a thought—thinkable—that is true, not what *makes* it true or “corresponds to it,” in any sense.) One form broadly truth conditional approaches to semantics can take is to construe the truth conditions of sentences as *intensions*: functions from indices to truth-values, thought of as the extensions of sentences. The worldly items that serve as arguments for such intension functions, and hence are assimilated as truth-makers or falsifiers of sentences, can be complete possible worlds or more fragmentary states or situations.

Thus far I have been concerned to emphasize the structural homologies between talk of paired truth-values true and false in a semantic metavocabulary and talk of paired practical attitudes of acceptance and rejection in a pragmatic metavocabulary. I think there are structurally similar reasons to look for something in addition to sentences on the side not of what *is* true/false or accepted/rejected, but what in some sense *makes* (what is expressed by) declarative sentences true or false and *makes* attitudes of acceptance or rejection appropriate or correct. There is a specifically semantic sense of correctness of acceptance/rejection that can just be read off of evaluations of truth and falsity, via the principle that accepting is taking-true and rejecting is taking-false. Acceptance is *semantically* correct just in case what is taken-true *is* true, and rejection is semantically correct just in case what is taken-false *is* false. And this non-causal, non-processual sense of “makes correct” can be expressed nonmetaphorically by a set of subjunctive conditionals: if the claim *were* true (its truth-conditions *were* satisfied), then acceptance *would be* semantically correct.

But there is another sense of “correctness” that applies to discursive practices, which confer on some acts and attitudes the practical significance of accepting or rejecting claimables

(truth-evaluables) expressed by (what thereby count syntactically as) declarative sentences. It concerns *justification*, rather than truth. This is a still-normative notion of correctness that is neither equivalent to the semantic one nor reducible to mere matter-of-factual identification of the psychological causal antecedents of actual adoption of practical attitudes of acceptance and rejection: what makes practitioners adopt in the sense of *causing* them to adopt those attitudes or the processes by which they do. It is a matter, rather, of the *reasons* interlocutors have for the attitudes of acceptance and rejection they adopt and express (in a pragmatic, rather than semantic sense) in their acts of affirmation and denial.

Appeals to the reasons to accept and reasons to reject various claimables that are available to discursive practitioners can play explanatory-explicative roles that are analogous in some important ways to the roles played by appeals to intensions and truth conditions. It is clear that there will be important structural disanalogies as well. For the reasons available to an interlocutor might justify neither acceptance nor rejection of some claimable, and one might have both some reasons to accept and some reasons to reject some (other) claimable.

I began by pointing to an analogy between the traditional semantic extensions of declarative sentences, the two basic truth-values, on the one hand, and the two fundamental pragmatic acts and attitudes articulating the use of such sentences, acceptance and rejection (practically *taking-true* and *taking-false*), on the other. I have now raised the possibility of extending that analogy from the extensional to the intensional, by looking at the *reasons* interlocutors can have to adopt practical attitudes of acceptance and rejection. The first step on the path to elaborating pragmatics by analogy to the way semantic intensions are elaborated from merely extensional truth-values is firmly to distinguish practices of reasoning from reason *relations*. This is to distinguish reasons *to* do something—adopt an attitude or perform an act, accept or reject—from reasons *for* and *against* claimables (what is expressed by declarative sentences in the sense of what can be accepted or rejected).

Gilbert Harman vividly illuminates the required distinction as part of his argument for the initially shocking claim that there are no such thing as rules of deductive reasoning. Paraphrased in the idiom I have been using here, he argues that if there were, presumably a paradigm would be the rule that if you accept p and accept $p \rightarrow q$, then you have decisive, deductively good reasons to accept q . So in those circumstances, you ought to do so. But, he points out, that would be a terrible rule. You might have much better reasons to reject q than you have to accept p or $p \rightarrow q$. In that case, you should reject one of them. Acceptance of conditionals can be exploited in reasoning either by *modus ponens*, or, equally validly, contrapositively by *modus tollens*.

What deductive logic directly supplies is reason *relations*. They tell us that some claimables provide reasons for and against others. Those relations are indeed relevant to

practices of reasoning, but only indirectly. They constrain but do not direct the drawing of conclusions, the adoption of some doxastic attitudes as justified by the adoption of others. The fact that p and $p \rightarrow q$ stand in the relation of *implication* to q tells us that one *ought not* accept p and $p \rightarrow q$ and reject q . (We will be normatively “out of bounds” if we do.) But it does not tell us what to *do* should we find ourselves with those attitudes—which one or more of them we should change. The fact that p and $\neg p$ stand in the relation of *incompatibility* (here, formal incompatibility: inconsistency) tells us that we *ought not* to accept both. But again, it does not tell us what to do in such a situation—which attitude we should give up.

Harman is concerned with the role of specifically *logical* reason relations. But claimables (acceptable/rejectables) expressed by sentences that contain no logical vocabulary also stand to one another in relations of implication and incompatibility. (Following Wilfrid Sellars, we may call these *material* reason relations.) “Plane figure A is square” implies “Plane figure A is polygonal.” So one ought not to accept the first and reject the second. “Plane figure A is square” is incompatible with “Plane figure A is circular.” So one ought not accept both. Nested kinds support implications: “Coda is a dog” implies “Coda is a mammal,” which implies “Coda is a vertebrate,” and so on. And contrary properties support incompatibilities: “Coda is a mammal” and “Coda is a reptile,” “Monochromatic figure A is green,” and “Monochromatic figure A is red,” and so on.

Harman argued that from relations of implication and incompatibility we cannot directly read off which acceptances and rejections give us reasons to accept and reject other claimables. We might, however, hope to be able to go the other way around. This would be to read off relations of implication and incompatibility between claimables from how some attitudes provide reasons to adopt other attitudes, via the non-determinative normative constraints that the reason relations impose on rationally admissible changes of attitudes. Pursuing this order of explanation would take us from practices of giving and assessing reasons to accept and reject the claimables expressed by declarative sentences to reason relations of implication and incompatibility that those claimables stand in to one another. The aspiration would be then to understand those claimable contents—what can be accepted or rejected in reasoning—in terms of the reason relations they stand in to one another.

One way to begin is to take it that

- i) If accepting A functions practically as a reason *to accept* B, then A provides a reason *for* B, and
- ii) If accepting A functions practically as a reason *to reject* B then A provides a reason *against* B.

As the next step, we can then think of *implication* relations as codifying reasons *for* and *incompatibility* relations as codifying reasons *against*.

The idea is to start with practices of reasoning, in the sense of practices of giving reasons that entitle one to acts or attitudes of accepting and rejecting claims. We can think of a dialogical situation, where those who accept or reject a claim can be challenged to defend that attitude, to *justify* it by offering *reasons* to accept or reject it. These practices of asking for and offering reasons *to* do something, to accept or reject a claim (claimable), must respect reason relations among claimables according to which some of them provide reasons *for* and reasons *against* others. These we understand as relations of material implication and incompatibility. What stand in *these* relations are not acts or attitudes, but claimable contents: what one can accept or reject (whether reasonably or not, depending on what reasons *to* adopt those attitudes one can offer).

I gestured earlier at theoretical reasons to think that reasoning practices must include the possibility of offering and assessing reasons to accept. This is to rule out the ultimate intelligibility of purely *skeptical* reasoning practices: practices that permit the adoption and justification only of attitudes of *rejection*. (The challenge of making sense of first-person practical reasoning, practical deliberation, was offered as a suggestive case in point.) In the present context, any such considerations provide reasons to think that reason *relations* must include *implications*, which codify reasons *for*. Reasons to think that anything intelligible as reasoning practices must also include the possibility of offering and assessing reasons to *reject* are not far to seek. These would deny the ultimate intelligibility of purely *dogmatic* reasoning practices: practices that permit the adoption and justification only of attitudes of *acceptance*. In the present context, any such considerations provide reasons to think that reason *relations* must include *incompatibilities*, which codify reasons *against*.

In a wonderful essay called “Why ‘Not’?”, Huw Price considers the practical deficiencies of what I am calling “purely dogmatic” reason-giving practices.³ He imagines “ideological positivists,” who do not have a way of denying or rejecting a claim. They accordingly lack any practical acknowledgment of the *incompatibility* of two claims. (It will follow that in their logic they have no way of *negating* a claim—hence the issue of his title.) He illustrates why such practices wouldn’t work with a nice dialogue:

Me: ‘Fred is in the kitchen.’ (Sets off for kitchen.)

You: ‘Wait! Fred is in the garden.’

Me: ‘I see. But he is in the kitchen, so I’ll go there.’ (Sets off.)

You: ‘You lack understanding. The kitchen is Fred-free.’

Me: ‘Is it really? But Fred’s in it, and that’s the important thing.’

(Leaves for kitchen.)

³ Huw Price “Why ‘Not’?” *Mind*, New Series, Vol. 99, No. 394 (Apr., 1990), pp. 221-238. Published by: Oxford University Press on behalf of the Mind Association. Stable URL: <https://www.jstor.org/stable/2254961>. Dialogue quoted is from p. 224.

Unless the claims we accept can *exclude* some other acceptances, they can't guide our actions. The essential conceptual starting-point of Shannon information theory—well upstream of the issue of how to *quantify* information—is the idea that if a message does not exclude some alternatives that were previously open, it conveys no information at all. We would learn *nothing* practically from finding out that there are reasons *for* someone to accept a claim—say, “Fred is in the garden,”—unless those same considerations can serve also as reasons *against* accepting some other claims—“Fred is in the kitchen,”—which accordingly count as *incompatible* with the original claim. That means that the very same claim that is a reason *for* one commitment must also be a reason *against* some others. Not only must it be possible to accept or to reject any claimable, in addition, adopting either of those attitudes towards a claimable must be able to serve *both* as a reason to *accept* some further claimables (a reason *for* them), *and* as a reason to *reject* some other claimables (a reason *against* them). What can be accepted or rejected must stand both in relations of implication and in relations of incompatibility. A discursive practice cannot be intelligible as articulated by one sort of reason relation unless it is intelligible as articulated by the other as well.

I have argued that what can be accepted must be capable of being rejected, and what can be rejected must be capable of being accepted, and that what can serve as a *reason to* accept some acceptable/rejectable must be capable of serving as a reason *to* reject other acceptables/rejectables, and *vice versa*. I have accepted Harman's argument that we should distinguish between norms governing conditional practical attitudes of acceptance/rejection and the reason *relations* that constrain, but do not determine those norms. As a result, I have argued, we must understand what can be accepted or rejected as standing in *both* sorts of reason relations: *implications*, codifying reason-*for* relations, and *incompatibilities*, codifying reason-*against* relations. All of this is a way of implementing the strategy of appealing to practical attitudes of accepting and rejecting what is expressed by sentences in order to understand the acceptable/rejectable *contents* expressed by declarative sentences by looking first to *reasons* other attitudes provide to accept or reject. The connection permitting this transition is supplied by the principle that a reason *to* accept (adopt that attitude) is governed by a relation between reasons *for* the content or object of that attitude (what is accepted), and that a reason *to* reject (adopt that attitude) is governed by a relation between reasons *against* the content or object of that attitude (what is rejected).

In this way we move from the idea of practical attitudes providing reasons *to do* something (adopt other attitudes) to relations of implication and incompatibility (reasons for and against) relating what can now be understood as what those attitudes are attitudes towards. It is an explanatory advance from *pragmatics*, studying what one is doing in adopting discursive attitudes, to *semantics*, studying the contents of those attitudes. Those contents are now thought of as nodes in a network of relations of implication and incompatibility.

To say that a set Γ of acceptables/rejectables *implies* acceptable/rejectable A, we can write “ $\Gamma \sim A$.” Use of the “snake turnstile” rather than the more familiar double turnstile \models of semantic consequence or the single turnstile \vdash of derivability reminds us that we are expressing *material* implications, not *logical* implications. To say that a set Γ of acceptables/rejectables is *incompatible* with acceptable/rejectable A, we can write “ $\Gamma \# A$.” (For my purposes here it suffices to stick to the more familiar and intuitive single-succedent notation. I’ll have something to say later about the multisuccedent analogues.)

So far my discussion of the attitudes of acceptance and rejection and (so) of reasons-for and reasons-against in the form of implications and incompatibilities has been reasonably even-handed. The picture has been symmetrical. There is a substantial structural asymmetry between the two kinds of reason relation however. Implication is not in general a symmetric relation. If, possibly in the context of Γ , A implies B, it does not at all follow that in the same context B implies A. $\Gamma, A \sim B$ does not entail $\Gamma, B \sim A$. “Pedro is a donkey” implies “Pedro is a mammal,” but not *vice versa*. By contrast, incompatibility is *de jure* symmetric. $\Gamma, A \# B$ *does* entail $\Gamma, B \# A$. “Oscar is an octopus” is incompatible with (a decisive reason against, a dispository reason to reject) “Oscar is a mammal,” and “Oscar is a mammal” is incompatible with (a decisive reason against, a dispository reason to reject) “Oscar is an octopus.”

We can ask: is there some reason why one reason relation is symmetric and the other not? Must it be so? What defect would a discursive practice have if it did not exhibit this structural asymmetry between the two kinds of reason relation? And if it there must be such a structural asymmetry, is it necessary that it be reasons *for* (codified in implications) that are nonsymmetric and reasons *against* (codified in incompatibilities) that are symmetric? Is it so much as intelligible that for some discursive practices it should be the other way around?

In *Making It Explicit* I explicated discursive practice in a normative pragmatic metavocabulary of *commitment* and *entitlement* to commitments (a framework whose rationale I shall return to in the next section). In those terms, I understood what it is for A to be incompatible with B as commitment to A precluding entitlement to B. (It is not that it is *impossible* to commit oneself to B by asserting it even though one is already committed to A. It is just that if one does, one has foregone the possibility of having the status of entitlement to B. Other interlocutors who are aware of the incompatibility will not treat one as justified in claiming B, as having an entitlement to it that might be inherited by others testimonially, for instance.) I think there is much to recommend such an understanding. But it is at least not obvious on such an analysis why the fact that commitment to A precludes entitlement to B should entail that commitment to B precludes entitlement to A. It seems possible that these could come apart, that we should keep separate sets of books on whether commitment to A

precludes entitlement to B and whether commitment to B precludes entitlement to A. But we don't find examples like this in the wild. Why not?

It might seem that this issue is wholly an artefact of the pragmatics-first order of explanation pursued here, that it is spurious, so that its arising at all counts substantially against that way of proceeding. For the traditional semantic theorist who treats logical vocabulary as available at the explanatory ground-level has a quick answer to the question about the rationale of the symmetry of incompatibility. To say that p and q are incompatible is to say that $p \rightarrow \neg q$. That is truth-functionally equivalent to $\neg(p \& q)$ and to $\neg p \vee \neg q$, which are symmetric because conjunction and disjunction are. So from a truth-functional perspective, $p \rightarrow \neg q$ iff $q \rightarrow \neg p$. That is just the symmetry of incompatibility, which we can codify in the Sheffer stroke.

Doesn't that analysis answer the question I asked and settle the issue I raised about the symmetry of incompatibility? It certainly shows that the symmetry of incompatibility is built deeply into semantic bivalence. So *if* the truth-functional connectives are the right way to codify and express material incompatibility, then it is symmetric. But they are *only if* material incompatibility, like logical inconsistency, is symmetric. It does in fact seem to be, which is a point in favor of the bivalent logical semantic analysis. But that analysis presupposes the symmetry of incompatibility, rather than explaining it—at least from the pragmatic point of view. We could say “So much the worse for the pragmatics-first order of explanation.” But I think if we dig deeper we will find a satisfying normative pragmatic explanation for the necessity of the symmetry of incompatibility. The invisibility of the question from the bivalent semantic logical point of view then must count against that order of explanation.

One mark of that invisibility, or better, technique for achieving it, is notational. In Gentzen-style sequent calculus formulations of reason relations, there is no separate sign for the relation of incompatibility. Incompatibilities are notationally assimilated to implications. To say that Γ is incompatible with A (what I am expressing by “ $\Gamma \# A$ ”), in a system of single-succedent sequents we write something like “ $\Gamma, A \mid \sim \perp$.” This attributes a property, call it “incoherence” to everything on the left of the turnstile, and says which property it is by using the perp sign for absurdity. Incoherence is expressed as implication of the absurd. This notation obviously builds in the symmetry of incompatibility, since “ $\Gamma, A, B \mid \sim \perp$ ” says *both* that, in the context of Γ , B is incompatible with A *and* that, in the context of Γ , A is incompatible with B (and similarly for any element of Γ). In multisuccedent sequent calculi the same effect is achieved—even less transparently—by foregoing the special sign for absurdity and having an empty right-hand side. It is a measure of the success of this notational sleight-of-hand in distracting from and hiding the commitment to the symmetry of incompatibility that Gentzen, who basically agrees with Tarski about the structural properties of logical consequence relations (even though he uses lists of sentences where Tarski uses sets), does not treat the symmetry of incompatibility as one of his

structural principles of reason relations. He does not enumerate it alongside monotonicity and transitivity (his “Thinning” and “Cut”).

III. Bilateral Pragmatic Renderings of Reason Relations

I began by distinguishing a *semantic* order of explanation, whose key conception is a distinction between two truth-values, true and false, with a *pragmatic* order of explanation, whose key conception is a distinction between two acts or practical attitudes, acceptance and rejection. In both cases there is a need to postulate something beyond declarative sentences, something that is in some sense *expressed by* such sentences, that is what is in the first instance true or false, accepted or rejected: the bearers of the truth-values, the objects of the doxastic attitudes. A traditional semantic approach is to understand propositional contents as intensions: functions from a set of arguments to truth-values. I have been sketching an alternate approach, within the pragmatic order of explanation. It looks instead to *reasons* interlocutors can have to adopt the basic attitudes of doxastic acceptance and rejection.

The idea here is that the practices within which acts and dispositions to act acquire the practical significance of doxastic acceptance and rejection are essentially, and not just accidentally, practices in which it is appropriate to ask for *reasons entitling* one to or *justifying* one in accepting or rejecting. We may use the term “claimable” as shorthand for “acceptable or rejectable”—taking it that it is the *same* kind of thing that can be accepted (taken-true) that can be rejected (taken-false) and *vice versa*, just as it is the same kind of thing that can be true that can be false. (That it is the same kind of thing that can be both true or false, on the semantic side, and accepted or rejected, on the pragmatic side, is a further, potentially more controversial matter.) Then the thought is that practices of accepting and rejecting are always also practices that include defending and challenging claimables, by offering reasons to accept or to reject them. Of course we can think of species of acceptance and rejection that are not like this. Voters might express their acceptance or rejection of a proposal by ticking one or the other of two boxes, with the issue of their *reasons* for doing so not being at issue (at any rate at for this purely voluntarist phase of the political process). But it seems clear that in order to understand *what* those voters are accepting/rejecting, we need to look beyond this aspect of their practices. The suggestion is that for this purpose it might be sufficient (and perhaps even necessary) to look to what would count as reasons to accept or reject the proposals they are voting on (and perhaps, what further acceptances and rejections the attitudes expressed by voting one way or another would provide reasons for).

I invoked an argument made by Gil Harman to point out the need to distinguish norm-governed practices of reasoning—*doing* something, taking up stances or positions and making moves to and from them—from reason *relations* among the claimables that define those positions (nodes in the network of relations) and license those moves. The practices consist of making claims (asserting and denying claimables), and defending and challenging them by

offering reasons to accept or reject them. Offering such a reason is just accepting or rejecting some other claimables (acceptables/rejectables), which stand in the right reason-relations to the one being defended or challenged. Those reason-relations are *implication* and *incompatibility*, relations in which some claimables serve as reasons for or against other claimables, respectively. The connection between offering reasons *to* accept or reject and the relations of being a reason for or a reason against a claimable consists, to begin with, in reasons *to* accept being reasons *for* the claimable accepted (implying it) and reasons *to* reject being reasons *against* the claimable being rejected (being incompatible with it). The culmination of this pragmatic order of explanation would be an account of claimables—the propositional contents that are expressed by declarative sentences and can be accepted or rejected—in terms of the role the sentences that express them play in relations of implication and incompatibility. What is envisaged is a pragmatic route to a kind of semantics, in the sense of an account of propositional contents (claimables).

Before considering that final step, in order properly to understand relations of implication and incompatibility, what is expressed by the snake and hash turnstiles in metalinguistic statements of the form “ $\Gamma \sim A$ ” and “ $\Gamma \# A$ ”, we must look more closely at the reasoning practices that they codify. For in the context of the pragmatic order of explanation being considered, the only grip we have on these relations is the role they play in practices of defending and challenging claims, by giving reasons to accept or reject them. Here there are two main points that I would like to argue for. The first is that to be intelligible as practices of reasoning, in the sense of accepting and rejecting claimables and defending and challenging those stances with reasons for and against them, the participants in such practices must be understood as keeping track of two different normative statuses: the kind of *commitment* one undertakes or acknowledges in accepting or rejecting a claimable by asserting or denying a sentence expressing it, and the sort of *entitlement* to that status or practical attitude that is at issue when *reasons* are offered for or against it. The second is that there is an important dimension along which these two flavors of normative status have quite different structures. The basis on which commitments are attributed is atomistic, while the basis on which entitlements are assessed is holistic.

As to the first point, we can begin with the observation that accepting or rejecting a claimable, paradigmatically by asserting or denying it, is taking a stand on it, adopting a stance towards it. It is committing oneself with respect to it, in the way one would by saying “Yea” or “Nay” to it in response to a suitable yes/no question. On the side of uptake, what some other practitioner needs to be able practically to discriminate in order to count as understanding the speech act is *that* the speaker has committed herself (performed a committive act, expressed a doxastic attitude), *how* she has committed herself (which kind of attitude she has adopted and expressed: acceptance or rejection, a positive or a negative commitment), and *to what* she has committed herself (toward which claimable she has adopted a doxastic attitude by asserting or denying the declarative sentence she uttered).

What difference does it make whether an interlocutor can offer reasons to accept what he has accepted or to reject what he has rejected? The doxastic *commitments* involved, the stances taken up, the attitudes adopted, are the same either way. But it is also an integral feature of doxastic commitments that one's *entitlement* to those (perhaps loosely undertaken) commitments is always potentially at issue. For in taking up a doxastic stance one renders oneself liable to demands for justification, for exhibition of reasons to accept or reject the claim one has accepted or rejected. One's liberty to commit oneself, to adopt that attitude and acquire that status, is not license to do so. Reasons matter because other practitioners must distinguish between the acceptances and rejections the speaker in question is *entitled* to, in virtue of having *reasons* to adopt those attitudes, and those the speaker is *not* entitled to, because unable to defend those commitments by offering reasons when suitably challenged to do so. It follows that for each interlocutor there must be not only a difference between the attitudes (commitments) he has adopted and those he has not, but also, within those he has adopted, between those he is entitled to or justified in, has rational credentials for, and those that are *mere* commitments, bare of such accompanying entitlements. In *Making It Explicit* I argue that what turns practically on one's entitlement or justification is the testimonial authority of one's act: its capacity to license others to adopt a corresponding attitude. The essential point is that in addition to the *committive* dimension of assertional practice, there is the *critical* dimension: the aspect of the practice in which the rational propriety of those commitments, their justificatory status, is assessed. Apart from this critical dimension, the notion of *reasons* gets no grip. It gets its grip from those keeping deontic score on their fellow discursive practitioners, who treat a failure to satisfy the justificatory responsibility implicit in undertaking a doxastic commitment as undercutting the interpersonal authority such a commitment otherwise could exercise.

The second structural observation is that entitlements are fragile in a way that commitments are not. They are vulnerable to being undercut by incompatible collateral commitments. The basic phenomenon here is twofold. First, it is not impossible for someone to be committed both to accept and to reject the same claimable. But second, one cannot then count as *entitled* to those contrary commitments. For each commitment provides a decisive reason against the other. The contrary commitments might have arisen through affirmation and denial of the same sentence—or, more commonly, when one is a (possibly unacknowledged) consequence of other attitudes one has self-consciously adopted. The mutual repulsion between the commitments implicit in attitudes of acceptance and rejection takes place at the normative level of rational entitlements to those commitments.

The origin and paradigm of the incompatibility of commitments undercutting their entitlements is the normative collision that occurs when one accepts and rejects the same claimable. But the phenomenon is not limited to that original case where contrary attitudes are adopted towards one and the same claimable. One treats the contents of two claimables as

incompatible just by taking it that commitment to one precludes entitlement to the other. I forfeit entitlement to my commitments if I both affirm and deny (accept and reject) that the plane figure is a circle. But I incur the same normative cost if I both accept that it is a circle and accept that it is a triangle. That is the practical normative significance of “A is circular” and “A is triangular” standing in the reason relation of material incompatibility (Aristotelian contrariety): commitment to one precludes entitlement to the other.

The fact that claimables stand to one another in the reason relation of incompatibility—the fact that commitment to one can preclude entitlement to the other—means that there is a structural asymmetry between the normative statuses of commitment and entitlement, which articulate essential dimensions of the practice of giving and asking for reasons, making claims and defending and challenging them. Knowing an interlocutor’s attitude toward a claimable, whether they accept or reject it, is sufficient to settle their commitment with respect to it. But to assess their *entitlement* to that commitment we have to consult all their other commitments. It is not enough that they can cite collateral commitments that provide good reasons for the commitment in question. It is necessary also that they not in addition have undertaken commitments that provide equally good reasons against it.

Ryan Simonelli has shown how we can assemble these conceptual raw materials so as to characterize both incompatibility and implication in terms of some sets of commitments precluding entitlement to others.⁴ In the normative pragmatic vocabulary put in play here, we can define the reason relations like this:

Implication (IMP):

$\Gamma \mid \sim A$ iff commitment to accept all of Γ precludes entitlement to *reject* A.

Incompatibility (INC):

$\Gamma \# A$ iff commitment to accept all of Γ precludes entitlement to *accept* A.

On this account, a reason *against* a *rejection* is an implication with that conclusion, since $\Gamma \mid \sim A$ says that commitment to all of Γ precludes entitlement to reject A. That is the same as a reason *for* an *acceptance*. Dually, an incompatibility $\Gamma \# A$ exhibits its premises Γ as providing both a reason *against* acceptance and (so) a reason *for* rejection.

Notice that these pragmatic readings respect Harman’s point. They do not directly dictate what inferences one draws. They address merely the rational cotenability of various attitudes. I have suggested that there is a useful sense in which the semantic contents of the claimables (acceptables/rejectables) expressed by declarative sentences can be understood as consisting in the roles they play in reason relations of these two kinds. These definitions show how semantic contents in this sense can be understood in purely pragmatic terms of commitments and

⁴ [ref. to forthcoming...?] In presentations to the “Research on Logical Expressivism” working group aimed at reconciling Restall-Ripley bilateralism with Smiley-Rumfitt bilateralism in something like the deontic pragmatic metavocabulary of *MIE*.

entitlements to accept and reject the claimables that stand to one another in relations of implication and incompatibility. We see here deep connections among the paired distinctions between acceptance/rejection, commitment/entitlement, and implication/incompatibility.

Although I have presented it for the more familiar single-succedent turnstile, that pragmatic definition of implication is recognizably a version of the bilateralist reading Greg Restall and David Ripley recommend in order to make sense of the multisuccedent relation of implication.⁵ Their account has the immediate benefit of demystifying multiple conclusion implications, which many have found hard to parse and motivate. (Why is the comma conjunctive on the left of the turnstile, when combining premises, and disjunctive on the right, when combining conclusions?) They recommend that we understand what is expressed when we write “ $\Gamma \mid \sim \Delta$ ” (to put it in the notation I am using here) for sets of sentences Γ and Δ as the claim that the normative position of anyone who is committed to accept all the sentences in Γ and reject all the sentences in Δ is “out of bounds.” This philosophically powerful pragmatic interpretation allows them to understand sequent calculi as consisting of rules that tell us that if some positions are out of bounds, then some others are also. It is then easy to see how logic, so construed, normatively *constrains* reasoners guided by the aim of remaining “in bounds,” without pretending to dictate unique answers to the question of what they should *do*, how in particular they should alter their commitments, in particular situations. Simonelli translates their normative pragmatic idiom into the vocabulary of commitment and entitlement familiar from *Making It Explicit*, and then shows how to give a parallel explicit treatment of incompatibility (expressed in Gentzen’s multisuccedent sequent calculi by sequents with empty right-hand sides) using the common structure of some commitments precluding entitlement to others.

Even if one acknowledges the felicity and fecundity of Restall and Ripley’s bilateral interpretation of the multisuccedent turnstile, one might still find it puzzling that implication relations should be thought of in terms of preclusion of entitlement. The standard way of thinking about implication (implicit in Tarski’s formal articulation of it and explicitly endorsed in *Making It Explicit*) seems rather to involve some commitments having other commitments as consequences. Is the account being put forward here committed to that idea just being wrong about implication relations? Is there really nothing to be made of that line of thought? Further, it seems that if Γ is incompatible with A then accepting it *commits* one to *reject* A. Here again, I think Simonelli shows us the way. He points out in effect that we can introduce a pragmatic sense of “implicit” according to the two principles:

⁵ G. Restall 2005. “Multiple Conclusions.” In *Logic, Methodology, and Philosophy of Science: Proceedings of the Twelfth International Congress*, edited by Petr Hájek, Luis Valdés- Villanueva, and Dag Westerståhl, 189– 205. London: College Publications. G Restall 2009b. “Truth Values and Proof Theory.” *Studia Logica* 92 (2): 241– 64. G. Restall 2013. “Assertion, Denial, and Non- classical Theories.” In Tanaka et al. 2013, 81– 100. D. Ripley 2013 “Paradoxes and Failures of Cut.” *Australasian Journal of Philosophy* 91 (1): 139– 64.

Pragmatically Implicit Acceptance (PIA): Any set of commitments that *precludes entitlement to reject A* thereby *implicitly* commits one to *accept A*.

Pragmatically Implicit Rejection (PIR): Any set of commitments that *precludes entitlement to accept A* thereby *implicitly* commits one to *reject A*.

PIA together with IMP entails that if $\Gamma \sim A$, then commitment to accept all of Γ *implicitly* commits one to accept A. PIR together with INC entails that if $\Gamma \# A$, then commitment to accept all of Γ *implicitly* commits one to reject A. In this way we can reconstruct what is right about the thought that implication is a matter of acceptance of some premises having commitment to acceptance of a conclusion as its consequence and incompatibility is a matter of acceptance of some premises having commitment to rejection of a conclusion as its consequence. The connection between those characterizations and the modified bilateralist normative pragmatic construal is provided by the implicit acceptance and implicit rejection principles PIA and PIR.

Why should we accept those principles? In what sense are commitments to accept and to reject implicit in preclusions of entitlement to reject and to accept, respectively? The claim of PIA is that if commitment to accept Γ precludes entitlement to reject A, then that same commitment to accept Γ implicitly commits one to accept A. One option, rejecting A, has been ruled out. One could remain agnostic, neither accepting nor rejecting. But that's not right. After all, one of the options has been ruled out. One cannot become entitled to reject A. The only option left standing, the only one available that one could potentially be entitled to is accepting A. By hypothesis, one has not yet explicitly done that. But that attitude of acceptance is implicit in the ruling out (as something one cannot be entitled to) of the only other option, in the sense that it is the only option left open. This is not the same as actually adopting the attitude, and that is what we mark by calling the commitment to accept "implicit," by contrast to the actual, explicit adoption of it. It seems clear both that this is an intelligible *pragmatic* sense of "implicit commitment to accept" and that calling it that is motivated by the rendering impermissible of the only other active option, rejection, and the consequent relative pointlessness of remaining uncommitted. In Section V, I discuss a *semantic* notion of implicitness, according to which the *explicit* content of a set of premises are the premises that are its elements, while its *implicit* content consists of what it *implies*.

With the distinction between commitments and entitlements on board, and the example of Restall's and Ripley's normative pragmatic bilateralism in mind, we are in a position to get a clearer view of the phenomenon that led Harman to distinguish reason relations such as implication from reasoning practices such as inferring in the sense of accepting or rejecting some conclusion on the basis of accepting a set of premises. Put in our terms, he points out in effect that the implication $\Gamma \sim A$ need not entitle one who is committed and entitled to accept all of Γ to

accept A, even though the implication implicitly commits him to it. For the interlocutor might have much better reason $\Delta \# A$ against A than Γ provides for it. If the interlocutor is also committed and entitled to accept all of Δ , then $\Gamma | \sim A$ precludes entitlement to reject A and $\Delta \# A$ precludes entitlement to accept it. Each explicitly precludes entitlement to the commitment that the other implicitly requires. The reason relations determine that one may not draw both conclusions, for one is never entitled both to accept and to reject A. That position is “out of bounds.” Commitments to which one otherwise would be entitled can have that entitlement undercut by collateral incompatible commitments. But the colliding implications and incompatibilities only preclude joint entitlements. They do not say what individual commitments you should accept or reject, what attitudes one should adopt or revise in response to that normative constraint, in the particular practical context of a determinate set of prior commitments. The important basic lesson Harman teaches about relations of implication and incompatibility merely constraining rather than dictating reasoning practices—how one ought to adopt further attitudes in the light of their rational relations to one’s antecedent attitudes—does not at all turn on specifically *contrapositive* forms of reasoning from implications, and can easily arise in their absence. (The stripped-down dialogic practices sketched in Section VII below, which forbid revision of commitments but track the fluctuating entitlements that result from their accumulation are a case in point.)

Formulating the issue in the normative pragmatic metavocabulary I am recommending also sheds some new light on the puzzle about the symmetry of incompatibility I raised earlier. I mentioned then that thinking just in terms of commitment and entitlement does not evidently provide a reason why the fact that commitment to A precludes entitlement to B should entail that commitment to B precludes entitlement to A. It seems intelligible that one of these relations should hold without the other. Why shouldn’t incompatibility be nonsymmetric, as implication is? If we look at the normative definitions IMP and INC we get the beginning of an answer to this question. On the side of implication, it is intelligible that commitment to accept A should preclude entitlement to reject B without its being the case that commitment to accept B should preclude entitlement to reject A. That is exactly how it is with “Pedro is a donkey,” and “Pedro is a mammal.” If you accept that he is a donkey, you are not entitled to deny that he is a mammal. But you can legitimately accept that he is a mammal and deny that he is a donkey, since he might be a capybara. The implication “Pedro is a donkey” $| \sim$ “Pedro is a mammal” rules out the position in which one accepts that he is a donkey and denies that he is a mammal. It says nothing about the legitimacy of switching the doxastic valence of those attitudes.

By contrast, we are to read “Pedro is a donkey” $\#$ “Pedro is a capybara” as saying that acceptance of Pedro being a donkey rules out entitlement to accept Pedro being a capybara. By contrast to the case of implication, the valence of the commitments in the case of incompatibility is the same. The Simonelli reading of implication in terms of preclusion of entitlement is recognizably a version or the Restall-Ripley reading of one’s position being “out of bounds” if

one both accepts the premises and rejects the conclusion. Applying and extending that model to the case of incompatibility understands INC as a version of what we would put in Restall-Ripley bilateralist terms as that the position in which one accepts both the premises and the conclusion is out of bounds. But since acceptance is involved in both cases, this is saying that accepting all of the elements of the *set* one gets by adding the conclusion to the premises is “out of bounds.” That is just what Gentzen represents by a sequent with an empty right-hand side. That marks the *set* as incoherent. Incompatibilities read off of that incoherent set will be *de jure* symmetric.

Now I do not take this argument to be decisive. If $A \# B$, then commitment to accept A precludes entitlement to accept B *whether or not one is entitled to accept A*. It does follow that one cannot be entitled to accept both A and B. But it does not follow from that fact that mere commitment to B precludes entitlement to A. On the Simonelli reading, commitment to A precludes entitlement to B. (In the general—but still single-succedent—case, that $\Gamma, A \# B$ iff $\Gamma, B \# A$.) The symmetry argument needs that one cannot be entitled to accept all of both the left-hand and the right-hand side of the $\#$ turnstile. The extension of the analogy with the Restall-Ripley bilateralist reading of the implication turnstile to a reading of the incompatibility turnstile is suggestive, but not coercive. It is probative, rather than dispositive.

We can use the considerations assembled here to complete the argument, making visible the nature of the pragmatic necessity that incompatibility be understood as a symmetric relation. It requires looking more closely at the dialogical pragmatic context in which those who give and ask for reasons entitling interlocutors to their commitments defend and challenge those commitments. The basic idea here, too, is due to Ryan Simonelli.⁶ The basic dialogic significance of showing that someone’s position is “out of bounds” or that they are not entitled to one of their commitments is presumably to force them to change those commitments. In the case we are addressing, this means withdrawing a commitment shown to be incompatible with others that interlocutor has undertaken. Nonsymmetric incompatibilities cannot serve this purpose. Suppose $A \# B$ but not $B \# A$: commitment to accept A precludes entitlement to accept B, but commitment to accept B does not preclude entitlement to accept A. Now consider an objection to a speaker S who is already committed to accept A and who then asserts B. It is pointed out that S’s commitment to accept A is incompatible with commitment to accepting B, so that S cannot be entitled to that commitment. S might repair his normatively “out of bounds” situation in response to this objection by withdrawing the commitment to accept B. Or, S could hold onto the acceptance of B and withdraw commitment to accepting A. If S does either of these things, S will be back normatively “in bounds” as far as this incompatibility is concerned. But if makes the repair by withdrawing commitment to accept A, in the case where A is incompatible with B but B is not incompatible with A, he can immediately reassert A, committing to accept it once again. Then it is not open to the objector to point to his acceptance of B as making this move illegitimate, a commitment to which S cannot be entitled. For that he is already committed to

⁶ [ref.?]

accept B is by hypothesis no objection to his acceptance of A. It does not preclude entitlement to that acceptance. The upshot is that nonsymmetric reason relations of incompatibility would be of no practical use in criticizing the commitments of any interlocutor. For any lack of entitlement they invoke can be repaired just by withdrawing the antecedent commitment and then endorsing it once again.

This argument is in some ways analogous to Dutch Book arguments in rational choice theories. It shows that those whose commitments are normatively criticized by invoking nonsymmetric incompatibility relations can immunize themselves from the effects of such criticism by the simple mechanism of withdrawing prior commitments shown to be incompatible with, and so rule out entitlement to, subsequent commitments, and then reinstating those same commitments. It would be a mistake to think that this argument turns on contingent details of the dialogic rules for challenging an interlocutor's claims by offering reasons against them by exhibiting their incompatibility with other commitments—which, accordingly, preclude entitlement to those claims. One might be thinking of practices in which one was not permitted to withdraw earlier commitments, but “lost” the dialogical game if convicted of incompatible commitments. But the point goes deeper than that. Entitlement-precluding incompatibilities must be assessable with respect to the whole set of an interlocutor's commitments. The question is whether any of them preclude entitlement to any of the rest. Nonsymmetric incompatibility relations would require distinguishing between two interlocutors who had exactly the same commitments, finding one “out of bounds” and the other “in bounds” normatively, depending on the order in which they had acquired those commitments. The requirement that incompatibilities be symmetric is the requirement that entitlements be assessable relative to the whole set of commitments whose entitlements are being assessed. That rules out what we might call “doxastic hysteresis”: the path-dependence of entitlements consequent upon nonsymmetric incompatibility relations.

Part Two: Semantics and Logic

IV. Two Routes to Semantics

I began by pointing out that parallel to the fundamental semantic bipolarity marked by the paired opposite truth-values, truth and falsity, there is a corresponding pragmatic bipolarity marked by the paired opposite practical doxastic attitudes of acceptance and rejection. I suggested that at least an important part of the theoretical work done on the side of semantics by invoking intensions, as functions from indices to truth-values, might be done on the side of pragmatics by looking at *reasons* to adopt practical attitudes of acceptance or rejection. According to this line of thought, the contents expressed by the declarative sentences used in overt assertions and denials are articulated by two sorts of reason relations among claimables (what can be accepted and rejected): implication and incompatibility. These are relations of being a reason *for* and being a reason *against*. There is a path to understanding these reason-relations pragmatically, in terms of what one is *doing* in offering reasons *to* do something—namely accept or reject—by means of the principle that a reason *to accept* a claimable is a reason *for* it (a claimable that implies it) and a reason *to reject* a claimable is a reason *against* it (a claimable that is incompatible with it).

I then argued that any pragmatic metavocabulary that is expressively adequate to characterize practices that are discursive in the sense that they include the offering of reasons for and against assertions and denials (practices of defending and challenging claims), must acknowledge two flavors of normative status practitioners can have with respect to claimables. These are the sort of doxastic *commitment* undertaken by adopting attitudes of acceptance or rejection and the sort of doxastic *entitlement* that is at issue when reasons are offered for and against those commitments. They articulate a kind of justificatory *responsibility* (to defend with reasons, upon suitable challenge) and a kind of *authority* (consequent upon successful justifying defense by reasons) respectively. Following Ryan Simonelli's development of Restall and Ripley's bilateralist normative pragmatic understanding of the meaning of the multisuccedent sequents studied in Gentzen-style proof theories yielded a reading of the implication $\Gamma \sim A$ as holding just in case commitment to accept all of the premises Γ precludes entitlement to *reject* the conclusion A , and a reading of the incompatibility $\Gamma \# A$ as holding just in case commitment to accept all of the premises Γ precludes entitlement to *accept* A .

Guiding this narrative has been the idea that once we have an understanding from the side of pragmatics of the fundamental pair of opposite-but-complementary reason relations, implication and incompatibility, it will be possible to use them to explicate the acceptables/rejectables expressed by declarative sentences. Understanding what can be doxastically accepted or rejected

in terms of the roles declarative sentences can play in reason relations of implication and incompatibility would provide a purely pragmatic explication of a fundamental semantic concept: the concept of the *contents* expressed by those declarative sentences. What I want to do next is to explain two contributions to this enterprise that are made by recent work by two other members of the ROLE working group, Ulf Hlobil and Dan Kaplan. Hlobil offers an illuminating perspective on the relation between a pragmatic story along the lines I have been telling here and the best contemporary work in formal semantics. Kaplan shows in detail how a proper semantic account of the contents expressed by declarative sentences can be elaborated from the role those sentences play in reason relations of implication and incompatibility.

One of the most sophisticated, flexible, and expressively powerful formal semantic understandings of conceptual content available today is Kit Fine's truth-maker semantics.⁷ It is built on a space of what he calls "states." We are invited to think of the states as facts or situations, but the notion is an adaptable one, sufficiently general to include whatever it is that we might think of as making declarative sentences true or false. A subset of the space of states is distinguished as the *possible* states. The only structure imposed on the state space is a partial ordering of part-hood: some states are parts of others. It is assumed that every subset of the space has a least upper bound. It can be thought of as the *fusion* of the elements of the subset: the unique whole of which they are all parts. The content or proposition expressed by a sentence A is then specified bilaterally, as a pair of sets of states: those "verifying" states that would make it true and those "falsifying" states that would make it false.

Like intensional semantics appealing to possible worlds, truth-maker semantics advances from the fundamental opposition of truth and falsity to a notion of content as truth conditions. It is more general in including also a notion of falsity conditions, which are not assumed in general to be uniformly computable from the truth conditions. Its basic notion of a state is more capacious than that of possible world. Possible worlds are included as special cases of states. For two states can be defined as *compatible* just in case their fusion is one of the states distinguished as *possible*. And a state can be understood as a possible world just in case it is a maximal possible state, in the sense of containing as parts every state compatible with it. Further flexibility is secured by not restricting the state space to *possible* states, but embedding those in a larger structure that includes multiple distinct *impossible* ones. In addition, the mereological structure of the state space provides expressive resources in the truth-maker semantic metavocabulary that have no analogue in classical possible worlds semantics. Finally, the bilateral conception of content, including falsifiers as well as verifiers and not assuming that either sort of semantic interpretant can straightforwardly be computed from the other, turns out to pay large expressive dividends.

⁷ Introduced in "A Theory of Truth-maker Content I: Conjunction, Disjunction, and Negation" *Journal of Philosophical Logic* (2017) 46:625-674.

The truth-maker semantic framework permits various definitions of the reason relations of implication and incompatibility. As just indicated, a state t counts as incompatible with a set S of states just in case the fusion of it with all the states in S is an impossible state. We can then say that $\Gamma \# A$ just in case any fusion of verifiers of all the members of Γ with any verifier of A is an impossible state. On the side of implication, there are a number of different notions of semantic consequence definable in the truth-maker setting, and Fine considers it a signal virtue of his approach that it can express and compare such a variety of senses of “follows from.” For instance, Γ verifier-entails A in case every state that verifies all the sentences of Γ verifies A .

Ulf Hlobil shows how the truth-maker framework allows the definition of a further notion of implication, which Fine does not consider.⁸ We can say that $\Gamma \sim \Delta$ iff any fusion of a state that verifies all the members of Γ with a state that falsifies all the members of Δ is an impossible state. He invites us to compare this semantic notion of multisuccedent implication with Restall and Ripley’s bilateral pragmatic notion. Recall that they understand $\Gamma \sim \Delta$ as saying that any position that includes accepting all of Γ and rejecting all of Δ is normatively incoherent or “out of bounds”—as we have read it: one cannot be entitled to such a constellation of commitments.

It is clear that these are isomorphic understandings of implication. The role played in the truth-maker semantic definition by verifiers and falsifiers of sentences is played in the bilateral pragmatic definition by practical attitudes of acceptance and rejection of sentences. And the role played in the truth-maker semantic definition by the impossibility of the state that results from fusing those verifiers and falsifiers is played in the bilateral pragmatic definition by the normative incoherence (or “out of bounds-ness”) of the position that results from concomitant commitment to those acceptances and rejections. The isomorphism extends to incompatibility as well as implication. In the single-succedent formulation, we can lay alongside the truth-maker semantic reading:

$\Gamma \# A \Leftrightarrow$ the state resulting from *fusion* of any *verifiers* of all the members of Γ with any *verifier* of A is an *impossible* state,

the normative pragmatic reading:

$\Gamma \# A \Leftrightarrow$ the position resulting from *concomitant commitment* to *accept* all of Γ and to *accept* A is normatively *incoherent* (“out of bounds”)—a constellation of commitments to which one cannot be entitled.

I believe that this isomorphism between the definitions of reason relations of implication and incompatibility in the bilateral semantic framework of verifiers and falsifiers and the bilateral pragmatic framework of acceptance and rejection is deep and revealing. To begin with, it shows how the connection between two paired truth values and two paired doxastic attitudes

⁸ Ulf Hlobil “The Laws of Thought and the Laws of Truth as Two Sides of One Coin” [ROLE: July 1, 2021]. [Update [ref.] as needed.]

expressed in the principles that accepting is taking-true and rejecting is taking-false is reflected, and can be further elaborated at the level of the reason relations of implication and incompatibility that articulate the contents that can *be* true/taken-true and false/taken-false. In particular, substantial new light is shed on what one must *do* to count thereby *as* adopting a practical attitude of taking some claimable to be true or false when those attitudes are situated in the wider context of practices of giving reasons for and against claimables that are constrained by reason relations of implication and incompatibility. The isomorphic relation between what is expressed by semantic metavocabularies of truth-makers and false-makers and what is expressed by pragmatic metavocabularies of bilateral commitments and preclusions of entitlement clarifies the relations between what one is *saying* and what one must be *doing* in order to say that in using the object language those semantic and pragmatic metavocabularies address. In practically acknowledging that commitment to accept some claimables precludes entitlement to reject some others and to accept still others, practitioners are, we can now see, *thereby* taking it that the fusion of verifiers of the premises and falsifiers (respectively, verifiers) of the conclusions are impossible states.

Alethic modal relations of possibility, impossibility, and necessity are part of the essential structure of the worldly states and situations that, according to the truth-maker semantic model, *make* claimables true or false, and so are what is represented and talked *of* or thought *about* by the use of declarative sentences. Deontic normative relations of commitment, entitlement, and preclusion of entitlement are part of the essential structure of discursive practical attitudes adoption of which, according to the pragmatics-first model, is what practitioners must *do* in order thereby to count as taking or treating what is expressed by declarative sentences *as* true or false, thereby representing the world as being some ways and not others by saying or thinking *that* things are thus-and-so. The very same reason relations of implication and incompatibility, which articulate the claimable contents expressed by declarative sentences, what can both *be* true or false and be practically *taken* to be true or false by affirming or denying them, can be construed *equally* and *isomorphically* both semantically, in alethic modal terms of *making* true or false, and pragmatically, in deontic normative terms of the practical doxastic attitudes of *taking* true or false (accepting or rejecting).

In *A Spirit of Trust* I attribute a view recognizably of this shape to Hegel, under the rubric “bimodal hylomorphic conceptual realism.” He emphasizes reason relations of material incompatibility (Aristotelian contrariety) over those of implication or material consequence—his notion of “determinate negation” over his notion of “mediation”—though both are always in play. As I read him, Hegel begins with the thought that ways the world can objectively be, facts, are determinate just insofar as they exclude and entail one another in a way properly expressed in alethic modal terms. That the coin is copper makes it *impossible* that it remain solid at 1100 degrees Celsius and *necessitates* its being an electrical conductor. By contrast, our subjective takings of the world to be some way, thoughts, are determinate just insofar as they exclude and

entail one another in a way properly expressed in deontic normative terms. As I've suggested here that we put this point, my commitment to the coin's being copper precludes *entitlement to accepting* that it would remain solid at 1100 degrees Celsius and precludes *entitlement to rejecting* that it is an electrical conductor. One and the same determinate conceptual content, that the coin is copper, can take two forms, an objective one in which it is understood as articulated by relations of exclusion and consequence construed in the alethic modal vocabulary proper to the expression of laws of nature, and a subjective one in which it is understood as articulated by relations of exclusion and consequence construed in the deontic normative vocabulary proper to the expression of discursive practices. That is why I use the term "bimodal hylomorphism." The view is properly denominated conceptual "realism" because the very same conceptual content to which we adopt attitudes in thought is understood as present, albeit in a different form, in the objective world thought about. The world is accordingly construed as essentially always already in a thinkable shape.

The isomorphism Hlobil has worked out between Restall and Ripley's normative pragmatic bilateral construal of implication and incompatibility relations and a version of Fine's truth-maker semantics is a colorable contemporary development of a thought cognate to the bimodal hylomorphic conceptual realism I attribute to Hegel. It suggests how something like this thought can be worked out in detail. For it maps onto one another a semantic idiom of great power and flexibility and a pragmatic idiom that has shown its substantial utility in understanding sequent calculi. Each has been used to characterize the fine structure of reason relations in actual applications to multifarious different object vocabularies.

When I introduced the idea of a pragmatics-first order of explanation, which would start with practices of accepting and rejecting and giving and asking for reasons entitling one to adopt those attitudes (so, challenging and defending doxastic commitments), I held out the prospect of a recognizably semantic understanding of the claimables that can be accepted or rejected (taken to be true or false) made available in terms of the reason relations of implication and incompatibility those claimables stand in to one another. We have seen how such reason relations can be understood in normative pragmatic terms of commitment and (preclusion of) entitlement, and how those very same reason relations can be reconstructed in paradigmatically semantic terms of worldly states or situations taken to make claimables true or false. But although the truth-maker semantics underwrites both a notion of the content expressed by declarative sentences and reason relations of implication and incompatibility that can also be understood in a normative pragmatic metavocabulary, it does not explain truth-evaluable content by appealing to those reason relations. Rather, it explains both in terms of modalized spaces of worldly states verifying and falsifying claimables. Striking as the isomorphism is that Hlobil points out between truth-maker semantic construals of implication and incompatibility and normative pragmatic construals of them, it does not amount to an explanation of claimable

content *by means of* reason relations. So it does not by itself count as redeeming the promissory note I issued on behalf of a pragmatics-first order of semantic explanation.

To do that we can look to the implicational phase-space semantics (IPSS) developed by Dan Kaplan, a Pittsburgh member of our ROLE logic working group. It implements precisely what I have been promising: an understanding of what is expressed by declarative sentences in terms of the role those sentences play in reason relations of implication and incompatibility. In so doing it fulfills the defining aspiration of the philosophical tradition I call “semantic inferentialism.” It begins with what I regard as a remarkable conceptual innovation. Not only are the semantic interpretants it appeals to implications (and incompatibilities), so is what is interpreted. That is, the principal and original bearers of semantic significance are construed not as sentences, but as implications.

The points of an implicational phase space are *candidate implications* defined on a language L_0 thought of as a set of logically atomic sentences. The candidate implications are then all ordered pairs $\langle \Gamma, \Delta \rangle \in L_0 \times L_0$ of sets of sentences of the language. They are what we have been representing by statements formed using the snake turnstile “ $\Gamma \vdash \Delta$.” This is the sort of thing manipulated in proof-theoretic multisuccedent sequent calculi—and given normative pragmatic interpretations by Restall-Ripley bilateralism. As is usual in such calculi, incompatibility is represented by empty right-hand sides rather than by a distinctive sort of turnstile: “ $\Gamma, A \vdash \sim$ ” rather than “ $\Gamma \# A$ ”. (The empty right-hand side marks the incoherence of the set of premises that appears on the left-hand side of the turnstile. Note that this notational convenience builds in an assumption of the *symmetry* of incompatibility.) I call the points of the implicational phase space “candidate” implications because they do not represent *good* implications: just candidates for that status. The good implications, the ones that actually hold—intuitively, where the set on the right-hand side, taken disjunctively, is a genuine consequence of the set of premises on the left-hand side, taken conjunctively, are marked as members of a distinguished subspace \mathbf{I}_0 of good implications.

The third element of an implicational phase-space semantic model for a language L_0 —in addition to the space of candidate-implication points $L_0 \times L_0$ and the subspace of good implications \mathbf{I}_0 —is an operation \cup of *adjunction* of candidate implications. It is defined by:

Adjunction: $\langle \Gamma, \Delta \rangle \cup \langle \Phi, \Lambda \rangle =_{\text{df.}} \langle \Gamma \cup \Phi, \Delta \cup \Lambda \rangle$.

To adjoin two candidate implications one produces a third candidate implication by combining (in the sense of unioning) their premises and combining (in the sense of unioning) their conclusions. With the minimal candidate implication $\langle \emptyset, \emptyset \rangle$ playing the role of an identity element, adjunction is a commutative monoid on the space $L_0 \times L_0$. (More on that later.)

Each candidate implication can now be assigned, as its semantic interpretant, the set of candidate implications whose adjunctions with it yield good implications, implications in the distinguished set \mathbf{I}_0 .

\mathcal{V} -sets: $\forall x \in L_0 \times L_0 \quad x^\mathcal{V} =_{df.} \{y \in L_0 \times L_0 : x \cup y \in \mathbf{I}_0\}$.⁹

The \mathcal{V} -set (pronounced “vee set”) of a candidate implication $\langle \Gamma, \Delta \rangle$ is what you need to add (adjoin) to it to get a *good* implication. If $\langle \Gamma, \Delta \rangle$ is already a good implication (if it is in \mathbf{I}_0) that fact will be marked by the fact that the minimal candidate implication $\langle \emptyset, \emptyset \rangle$ will be in its \mathcal{V} -set. If $\langle \Gamma, \Delta \rangle$ is a good implication, its \mathcal{V} -set $\langle \Gamma, \Delta \rangle^\mathcal{V}$ is something like its *range of subjunctive robustness*. Focusing for simplicity on the premise-set Γ , the \mathcal{V} -set is telling us what further collateral premises we can add to it without infirming the implication: turning it from a good one to a bad one. If the hungry lioness sees a limping gazelle nearby, then she will pursue it. That implication would still be good even if the beetle on a distant tree climbs a bit further out on the branch as it is sitting on. But it would not be good if the lioness were suddenly struck by lightning. If the candidate implication is not a good one, its \mathcal{V} -set tells us what we would need to add (adjoin) to it to *make* it a good one. Intuitively, the \mathcal{V} -sets play a role with respect to implications that is analogous to the role played by truth conditions with respect to sentences. They both specify what it would take for one to be semantically good—in the (different) ways implications and sentences can be semantically good.

At a second, separate stage, this semantic interpretation *of* (sets of) implications *by* sets of implications can then be extended to specify the semantic roles played by *sentences in* implications (and incompatibilities), rather than just of the implications themselves. In this implications-first inferentialist setting, a sentence A can be represented for semantic purposes by a pair of implications: $\langle \langle A, \emptyset \rangle, \langle \emptyset, A \rangle \rangle$. The semantic content expressed by the sentence—in the sense of its role in reason relations of implication and incompatibility—can then be represented by the $\langle \emptyset, A \rangle^\mathcal{V}$ these paired implications. $\langle A, \emptyset \rangle^\mathcal{V}$ determines the set of all the good implications in which A figures as a premise. $\langle \emptyset, A \rangle^\mathcal{V}$ determines the set of all the good implications in which A figures as a conclusion. For each tells us what additions to the bare skeletons of $\langle A, \emptyset \rangle$ and $\langle \emptyset, A \rangle$ yield good implications. The nature of the adjunction operation in terms of which (together with the set of good implications \mathbf{I}_0) the \mathcal{V} -sets are defined ensures that A appears as a premise in every element of the set of good implications that results from adjoining elements of $\langle A, \emptyset \rangle^\mathcal{V}$ to $\langle A, \emptyset \rangle$, and as a conclusion in every element of the set of good implications that results from adjoining elements of $\langle \emptyset, A \rangle^\mathcal{V}$ to $\langle \emptyset, A \rangle$.

The claim is that broadly inferential roles, in the sense specified by pairs of premissory and conclusory \mathcal{V} -sets $\langle \langle A, \emptyset \rangle^\mathcal{V} \langle \emptyset, A \rangle^\mathcal{V} \rangle$ are a good representation of what one must grasp in

⁹ \mathcal{V} -sets can be computed for *sets* of implications by requiring that each element of the \mathcal{V} -set yield an element of \mathbf{I} when adjoined with *every* element of the set: $\forall X \subseteq L_0 \times L_0 \quad X^\mathcal{V} =_{df.} \{y \in L_0 \times L_0 : \forall x \in X [x \cup y \in \mathbf{I}_0]\}$.

order to understand what one is accepting or rejecting in undertaking doxastic commitments.¹⁰ For it is inferential roles in this sense that determine what is a reason for and against the claims to which one is committing oneself, and so what it would take to entitle oneself to those attitudes and the acts of affirmation and denial that overtly manifest them. For that reason, these are good semantic representations of the claimable contents expressed by declarative sentences. Of course the idea is not that in order to defend and challenge doxastic commitments we need to have fully mastered the intricacies of these inferential roles. It is that insofar as we do not, we do not know what we are committing ourselves to, do not fully understand what we are accepting or rejecting, or the reasons we give entitling us to do so.

A minimal criterion of adequacy for Kaplan's implicational phase-space semantics is that it can be shown to offer a tractable semantics for the logically complex sentences that result when we extend the logically atomic language L_0 by introducing sentential logical vocabulary according to a wide variety of sequent rules. Indeed, Kaplan proves soundness and completeness results using the implicational phase-space semantics for a number of such logics, including not only classical and intuitionistic logics, but also a wide variety of substructural (nonmonotonic, nontransitive, noncontractive...) logics—about which more later (in Section VI). This broadly inferentialist semantic account of the claimable (acceptable/rejectable) contents expressed by declarative sentences is what I had in mind when I initially raised the possibility that a pragmatics-first approach that understands reason relations of implication and incompatibility in normative terms of what one is *doing* in adopting doxastic practical attitudes of accepting and rejecting claims and challenging and defending entitlement to the resulting commitments by offering reasons for and against them could be built on, extended, and developed to provide an adequate semantics.

The title of this section is “Two Routes to Semantics.” I have described how Hlobil offers a way of understanding (his version of) reason relations, paradigmatically implication, in Fine's truth-maker semantics, in terms of an isomorphism with Restall and Ripley's bilateralist normative pragmatics. I have explained in general terms how Kaplan defines his implicational phase-space semantics directly in terms of implication (and incompatibility) relations, which we have seen can be understood in normative pragmatic terms of acceptance and rejection, commitments and (preclusions of) entitlement. I want to close this section by comparing and contrasting the reconstructions of reason relations of implication and incompatibility offered by these two semantic approaches: in terms of truth-makers and in terms of implications.

The first thing to appreciate is the strong formal analogies between the two frameworks. The modalized state spaces of truth-maker semantics are built on sets of “states” that are not further specified. The states making up these spaces could be pretty much anything—which of

¹⁰ For both conceptual and technical reasons, it turns out that it is best to use the closures of these \vee -sets under the \vee -function, which can be shown to reach a fixed point at $\langle\langle A, \emptyset \rangle^{\vee\vee} \langle \emptyset, A \rangle^{\vee\vee} \rangle$, but I suppress this complication.

course contributes greatly to the flexibility of the apparatus. Within the set S of states, a privileged subset of “good” ones, S^\diamond is distinguished—intuitively, by its alethic modal status as “possible.” Kaplan’s implicational phase spaces are sets of points that have more structure than Fine’s states. They are candidate implications, pairs of sets of sentences drawn from an antecedent prelogical language. Within this space $L_0 \times L_0$ of implications, a privileged subset of “good” ones, I_0 is distinguished—intuitively, by its normative status as codifying the proper implications, what really follows from what. The operations on states or candidate implications, fusion \sqcup and adjunction \cup (the one stipulated, the other defined in terms of the additional structure of the space of candidate implications defined on L_0) are algebraically both commutative monoids.¹¹ The semantic interpretants of sentences are in both cases bilateral: verifiers/falsifiers and premissory and conclusory \vee -sets respectively.

There is also a substantial formal difference between the two settings. In the truth-maker framework, the modalized state space with its fusion operation (or part-whole relation among states) is wholly distinct from the language it is used to interpret semantically. To get a semantic model, a third element is required: an interpretation function that maps sentences of the language onto pairs of sets of verifying and falsifying states. In the implicational phase-space framework, there is nothing corresponding to this extra element, connecting and mediating between the language and the space on which it is interpreted. The extra structure that the points of the implicational phase-space come with, their being candidate implications in the form of pairs of sets of sentences of the language, not only means that the monoidal operation of adjunction of candidate implications can be explicitly defined set-theoretically, as opposed to simply stipulated, as with fusion of states. Because the sentences themselves are already present in the space from which semantic interpretants are drawn, the \vee -function that semantically interprets first implications and then sentences can also be explicitly defined set-theoretically from the raw materials already present in the implicational phase-space itself. In this sense, the interpretation function connecting sentences to their semantic interpretants is *intrinsic* to the sentences as they figure in the space of implications. The sentences come already interpreted by the reason relations they stand in to one another, the roles they play in implications and incompatibilities. All the semantic framework does is draw that intrinsic interpretation out explicitly. Now whether this is a virtue or a vice, a benefit or a cost, will depend on collateral theoretical commitments. For one might see it as showing that the implicational phase-space framework is foolishly trying to do without relations to extralinguistic reality that are what make truth-maker semantics a genuine *semantics* in the first place. I am not going to argue about that. But I do want to assemble some further considerations that might bear on such a dispute.

¹¹ Both Fine’s truth-maker semantics on modalized state spaces and Kaplan’s implicational phase-space semantics use commutative monoids (the fusion/adjunction operation, together with a null space unit element) on spaces with distinguished subspaces (S^\diamond and I). This is an algebraic generalization of more familiar residuated lattices. In making this generalization, both are downstream from Girard’s phase-space semantics for linear logic.

For in spite of the substantial difference in the conceptions of semantic interpretation that animate the two different approaches, the fact that both take the mathematical form of commutative monoids plus distinguished subspaces means that their treatment of the crucial reason relations of implication and incompatibility share enough structure to be intertranslatable across the two settings. That is, we can specify exactly the same reason relations of implication and incompatibility while moving systematically between the modalized state spaces of truth-maker semantics and implicational phase-space semantics. Here's how.

For one direction: Beginning with a truth-maker model, one can define an implicational phase space that corresponds to it in the sense of defining exactly the same implications and incompatibilities. We are given a truth-maker model of a language L_0 , defined on a modalized state space $\langle S, S^\diamond, \sqcup \rangle$, which assigns to each sentence $A \in L_0$ a pair of sets of states $\langle v(A), f(A) \rangle$ understood as verifiers and falsifiers of that sentence. The points of the implicational phase space being defined are ordered pairs of sets of sentences of L_0 . These are the candidate implications. What corresponds to fusion, \sqcup , is adjunction: $\langle \Gamma, \Delta \rangle \sqcup \langle \Theta, \Psi \rangle = \langle \Gamma \cup \Theta, \Delta \cup \Psi \rangle$, as usually defined in implicational phase space semantics. It remains to compute \mathbf{I}_0 , the set of *good* implications. We do that using the consequence relation $\mathbf{H}_{\text{lobil}}$ defined to mimic the Restall-Ripley bilateral understanding of the multisuccedent turnstile:

$$\langle \Gamma, \Delta \rangle \in \mathbf{I}_0 \text{ iff } \forall s, t \in S [(\forall G \in \Gamma [s \in v(G)] \ \& \ \forall D \in \Delta [t \in f(D)]) \Rightarrow s \sqcup t \notin S^\diamond].$$

That is, $\langle \Gamma, \Delta \rangle$ is a good implication just in case the fusion of any state s that verifies all of Γ and any state t that falsifies all of Δ is an impossible state, in the truth-maker model. This construction obviously guarantees that exactly the same implications will hold in the implicational phase space, that is, be elements of \mathbf{I}_0 , as satisfy the $\mathbf{H}_{\text{lobil}}$ consequence relation in the truth-maker model.

As for incompatibilities, in the truth-maker setting, two *states* s and t are incompatible just in case their fusion is an impossible state. Two *sentences* A and B are incompatible just in case any fusion of a verifier of the one with a verifier of the other is an impossible state. More generally, a set Γ of sentences is *incoherent* in case any fusion of verifiers of all its elements is an impossible state. Given the definition of the set of good implications \mathbf{I}_0 just offered, this is equivalent to $\langle \Gamma, \emptyset \rangle \in \mathbf{I}_0$. The incompatibilities are represented in the implicational phase space semantics just by good implications with empty right-hand sides.

So there is a straightforward method for taking any truth-maker model defined on a modalized state space and defining from it an implicational phase space model that has exactly the same reason relations of implication and incompatibility.

For the other direction: Beginning with an implicational phase space, one can define a truth-maker model (an interpreted modalized state space) that corresponds to it in the sense of defining exactly the same implications and incompatibilities. We are given an implicational phase space

defined on a language L_0 , $\langle \mathcal{P}(L_0) \times \mathcal{P}(L_0), \mathbf{I}_0 \rangle$. The states will be candidate implications. $S = \mathcal{P}(L_0) \times \mathcal{P}(L_0)$. \sqcup is adjunction: $\langle \Gamma, \Delta \rangle \sqcup \langle \Theta, \Psi \rangle = \langle \Gamma \cup \Theta, \Delta \cup \Psi \rangle$. In the Hlobil truth-maker definition of consequence, the *good* implications correspond to *impossible* states. So the subset of *possible* states is defined by $S^\diamond = S - \mathbf{I}_0$. It remains to define the model function m , which assigns to each $A \in L_0$ a pair of subsets of S , $\langle v(A), f(A) \rangle$, where $v(A) \subseteq L_0$ and $f(A) \subseteq L_0$, such that:

$$\langle \Gamma, \Delta \rangle \in \mathbf{I}_0 \text{ iff } \forall s, t \in S [(\Gamma = \{G_1 \dots G_n\} \ \& \ g_1 \in v(G_1) \ \& \dots \ g_n \in v(G_n) \ \& \ s = g_1 \sqcup \dots \sqcup g_n \ \& \ \Delta = \{D_1 \dots D_n\} \ \& \ d_1 \in v(D_1) \ \& \dots \ d_n \in v(D_n) \ \& \ t = d_1 \sqcup \dots \sqcup d_n) \Rightarrow s \sqcup t \notin S^\diamond].$$

For various metatheoretic purposes, Fine employs “canonical” truth-making models, in which the verifier of a (logically atomic) sentence is just that sentence and the falsifier of that sentence is just the negation of that sentence. (His requirement that the fusion of any verifiers of A will be a verifier of A and the fusion of any falsifiers of A will also be a falsifier of A is then trivially satisfied, since there is only one.) We can combine that idea with Kaplan’s standard representation of the proposition expressed by A as the pair $\langle \langle A, \emptyset \rangle, \langle \emptyset, A \rangle \rangle$, and do without the formation of falsifying literals by appeal to negation by defining the verifiers of A by $v(A) = \langle A, \emptyset \rangle$ and the falsifiers of A by $f(A) = \langle \emptyset, A \rangle$. We want to implement Hlobil’s definition of implication, that an implication $\Gamma \mid \sim \Delta$ is good in the truth-maker setting just in case the fusion of any verifier of all of Γ and any falsifier of all of Δ is an impossible state. To do that, we need to say what it is for a state (defined in the implicational phase space, that is, a candidate implication) to “verify all of Γ ” and to “falsify all of Δ .” We can extend the single-sentence definitions as follows. If $\Gamma = \{G_1 \dots G_n\}$ and $\Delta = \{D_1 \dots D_m\}$:

$$v(\Gamma) = \langle \Gamma, \emptyset \rangle = \langle G_1, \emptyset \rangle \cup \dots \cup \langle G_n, \emptyset \rangle.$$

$$f(\Delta) = \langle \emptyset, \Delta \rangle = \langle \emptyset, D_1 \rangle \cup \dots \cup \langle \emptyset, D_m \rangle.$$

That is, the implication (standing in for a state) $\langle \Gamma, \emptyset \rangle$ counts as verifying all of Γ because it is the adjunction of the verifiers of each element of Γ . (In this “canonical” modalized state-space model, sets of sentences, like individual sentences, only have single states=implications as verifiers.) And similarly for falsifiers.

To show that this works, in the sense of yielding the same implications in the truth-maker model that are good in the original implicational phase space, we must show that

$$\langle \Gamma, \Delta \rangle \in \mathbf{I}_0 \text{ iff } \forall s, t \in S [(\forall G \in \Gamma [s \in v(G)] \ \& \ \forall D \in \Delta [t \in f(D)]) \Rightarrow s \sqcup t \notin S^\diamond].$$

To show the left-to-right direction \Rightarrow : If $\langle \Gamma, \Delta \rangle \in \mathbf{I}_0$ then $v(\Gamma) = \langle \Gamma, \emptyset \rangle$ and $f(\Delta) = \langle \emptyset, \Delta \rangle$. So $v(\Gamma) \sqcup f(\Delta) = \langle \Gamma, \Delta \rangle$. Since by hypothesis $\langle \Gamma, \Delta \rangle \in \mathbf{I}_0$, by the definition of S^\diamond as $S - \mathbf{I}_0$, it follows that $\langle \Gamma, \Delta \rangle \notin S^\diamond$, that is, that the state $\langle \Gamma, \Delta \rangle$ is an impossible state. It is the fusion of *the* verifier of Γ , $\langle \Gamma, \emptyset \rangle$ and *the* falsifier of Δ $\langle \emptyset, \Delta \rangle$ because it is the result of adjoining them.

To show the right-to-left direction \Leftarrow : If $\forall s, t \in S [(\Gamma = \{G_1 \dots G_n\} \ \& \ g_1 \in v(G_1) \ \& \dots \ g_n \in v(G_n) \ \& \ s = g_1 \sqcup \dots \sqcup g_n \ \& \ \Delta = \{D_1 \dots D_n\} \ \& \ d_1 \in v(D_1) \ \& \dots \ d_n \in v(D_n) \ \& \ t = d_1 \sqcup \dots \sqcup d_n) \Rightarrow s \sqcup t \notin S^\diamond]$, then $s = v(\Gamma)$ and $t = f(\Delta)$, so $v(\Gamma) \sqcup f(\Delta) = \langle \Gamma, \Delta \rangle \notin S^\diamond$. Since $S^\diamond = S - \mathbf{I}_0$ and $\langle \Gamma, \Delta \rangle \in S$, $\langle \Gamma, \Delta \rangle \in \mathbf{I}_0$.

As for incompatibility, we must show that A and B are truth-maker incompatible (Γ is truth-maker incoherent), that is, $\forall s, t \in S [s \in v(A) \ \& \ t \in v(B) \Rightarrow s \sqcup t \notin S^\diamond]$, (or more generally, $v(\Gamma) \notin S^\diamond$) iff $\langle \{A, B\}, \emptyset \rangle \in \mathbf{I}_0$ (or more generally, $\langle \Gamma, \emptyset \rangle \in \mathbf{I}_0$).

To show the left-to-right direction \Rightarrow : If $\forall s, t \in S [s \in v(A) \ \& \ t \in v(B) \Rightarrow s \sqcup t \notin S^\diamond]$, then since $v(A) = \langle A, \emptyset \rangle$ and $v(B) = \langle B, \emptyset \rangle$, and since \sqcup is adjunction, $s \sqcup t = \langle \{A\} \cup \{B\}, \emptyset \rangle = \langle \{A, B\}, \emptyset \rangle$. Since $\Rightarrow s \sqcup t \notin S^\diamond$, $s \sqcup t = \langle \{A, B\}, \emptyset \rangle \in \mathbf{I}_0$. This works for arbitrary iterations of \sqcup , which gives the more general Γ case.

To show the right-to-left direction \Leftarrow : If $\langle \{A, B\}, \emptyset \rangle \in \mathbf{I}_0$, then $\langle \{A\} \cup \{B\}, \emptyset \rangle \in \mathbf{I}_0$.

Since \sqcup is adjunction, $\langle A, \emptyset \rangle \sqcup \langle B, \emptyset \rangle \in \mathbf{I}_0$. But $v(A) = \langle A, \emptyset \rangle$ and $v(B) = \langle B, \emptyset \rangle$.

So $v(A) \sqcup v(B) \in \mathbf{I}_0$. Since $S^\diamond = S - \mathbf{I}_0$, $v(A) \sqcup v(B) \notin S^\diamond$. That is truth-maker incompatibility of A and B. This works for arbitrary iterations of \sqcup , which gives the more general Γ case.

So there is a straightforward uniform translation between Kaplan’s implicational phase-state semantics and Fine’s truth-maker semantics. Each truth-maker model on a language corresponds to an implicational phase-space defined on that same language, in the sense that they underwrite exactly the same reason relations of implication and incompatibility. The parallel extends to various structural constraints that can be placed on them—Fine’s Exclusivity, Downward Closure, and Exhaustivity conditions, which I’ll have more to say about further along.

I have been talking about how the “internal” consequence (and incompatibility) relations line up in the two settings. Looking somewhat further afield, the deep affinities between these two semantic approaches are also manifested in the way verifiers line up with premissory roles and falsifiers with conclusory roles, in the *external* consequence relations. (The internal relations cross the turnstile(s). The external ones remain on one side of the turnstile, looking at relations between the premissory sides of different sequents, or between the conclusory sides of different sequents. In substructural cases, the internal and external consequence relations can diverge.) Kaplan shows that K3 (the Strong Kleene three-valued logic) is the unilateral external logic of premissory roles in codifying the sense of consequence in which $A \mid =_p B$ just in case if in the internal logic $\Gamma, B \mid \sim \Delta$ then $\Gamma, A \mid \sim \Delta$ (A can replace B as a premise, saving the goodness of implications), and LP (Graham Priest’s “Logic of Paradox”) is the unilateral external logic of conclusory roles in codifying the sense of consequence in which $A \mid =_c B$ just in case if in the internal logic $\Gamma \mid \sim A, \Delta$ then $\Gamma \mid \sim B, \Delta$ (B can replace A as a conclusion, saving the goodness of implications). Hlobil shows that K3 is the unilateral external “logic of verifiers,” in the sense that K3 preserves compatibility with the verifiers of the premises (jointly) to the verifiers of the conclusions (separately). And Hlobil shows that LP is the unilateral “logic of falsifiers,” in the sense that LP preserves the compatibility potential of the falsifiers of the conclusions (jointly) to the falsifiers of the premises (separately). So the isomorphism between the reason relations

specified by the truth-maker semantics and those specified by the implicational phase-space semantics goes beyond the internal (bilateral) consequence relations all the way to the external (unilateral) consequence relations as well.

V. The Structure of Material Reason Relations

I have suggested a pragmatics-first order of explanation rooted in an understanding of the practical significance of adopting normative doxastic attitudes of acceptance and rejection. To be genuinely doxastic attitudes, accepting and rejecting must be the undertaking of commitments whose entitlements are always potentially at issue. Entitlements are vindicated, conferred, and demonstrated by offering *reasons* to adopt commitments to accept or reject. Those reasons *to* adopt normative doxastic attitudes of acceptance or rejection can in turn be understood as expressions of relations among acceptable/rejectable (claimable) contents expressed by declarative sentences. Reasons to accept are reasons for claimables, and reasons to reject are reasons against claimables. These are relations of implication and incompatibility. Thinking about the roles sentences can play with respect to these two sorts of reason relations, the positions they can occupy in a network of such relations, then offers a route to a semantic account of claimable contents. Beginning with a story about sayings as *doings*, we arrive—via considerations of practices of giving and asking for *reasons*, defending and challenging claims—at an account of what can be *said*.

In this pragmatics-first order of explanation, reason relations articulate the contents interlocutors can commit themselves to by accepting or rejecting them, and can, if all goes well, be entitled to. Two features of the fundamental structure of reason relations have emerged from this story. First, for deep pragmatic reasons, reason relations come in two flavors (are of two kinds): implication and incompatibility, corresponding to reasons for and reasons against. Second, implication is nontransitive, while incompatibility is, for deep pragmatic reasons, a symmetric relation. What else can we discover about the structure of the semantogenic reason relations? I have told the story so far without appeal to logic or logical vocabulary. (Even the appeal to negated literals that Fine makes in defining his canonical models was avoided or evaded by using Kaplan's premissory and conclusory roles in the place of the contrasting atomic and negated literals.) But this is one place where we might sensibly look for enlightenment from the side of logic.

For there is wide agreement about the structure of *logical* reason relations. In classic papers published in 1935-36 Alfred Tarski and Gerhard Gentzen founded rival traditions by introducing, respectively, model-theoretic and proof-theoretic metavocabularies for codifying relations of logical consequence.¹² In spite of their fundamentally different approaches, they impose essentially the same structural constraints (if we divide through by the fact that Tarski worked

¹² Alfred Tarski's 1936 classic "On the Concept of Logical Consequence" pp. 409-420 in *Logic, Semantics, Metamathematics* [Oxford University Press, 1956]. Gerhard Gentzen's 1935 "Investigations into Logical Deduction" [English translation: *American Philosophical Quarterly* Volume I, Number 4, October 1964, pp. 288-306].

with sets and Gentzen with lists). Tarski expresses them using Kuratowski’s axioms defining topological closure operators. Omitting irrelevances, they are¹³:

1. $\Gamma \subseteq \text{Con}(\Gamma)$.
2. $\text{Con}(\Gamma \cup \Delta) = \text{Con}(\Gamma) \cup \text{Con}(\Delta)$.
3. $\text{Con}(\text{Con}(\Gamma)) = \text{Con}(\Gamma)$.

The first says that the premise-set is contained in the consequences of those premises. I will call it “Containment” (CO). The third is the closure condition. In his sequent calculi, Gentzen gets the same effect as these by imposing:

- Reflexivity (RE) by requiring the leaves of his proof trees to have the form:

$$A \vdash A$$

- Monotonicity (MO)—his “Thinning”:

$$\frac{\Gamma \vdash A}{\Gamma, B \vdash A}$$

- Cumulative Transitivity (CT)—his “Cut” (in the single-succedent formulation):

$$\frac{\Gamma \vdash A \quad \Delta, A \vdash B}{\Gamma, \Delta \vdash B}$$

Tarski’s CO is a consequence of RE and MO, and MO and CT are equivalent to Tarski’s other two conditions. (For simplicity, I follow the tradition here in focusing on logical consequence or implication. As we will see, corresponding principles apply, and issues arise—concerning for instance monotonicity—for incompatibility relations as well.)

There are good reasons to impose these structural conditions on logical relations of consequence or implication. As to Containment, it does seem that among the things that follow from any set of premises we should count those premises themselves. Monotonicity says that if an implication is good, if some conclusion genuinely follows from a set of premises, that it remains good, the same conclusion still follows, if we add further premises. Transitivity says that we can use the conclusions we draw from some set of premises as further premises from which to draw further conclusions. These all seem like generally plausible principles of logical reasoning, and they are all satisfied in mathematical reasoning, which is the solidest and in many ways the best understood sort of reasoning we engage in. Further, sentential logical connectives can be introduced in particularly simple ways, and are particularly well-behaved in settings with this structure. There has been growing interest in substructural logics, especially nonmonotonic logics, and logicians addressing semantic paradoxes have found it useful to consider relaxing other conditions besides MO. Such enterprises, though, still have a somewhat suburban status, and are located by their relations to the downtown of the fully structural settings in which core logics, paradigmatically classical bivalent and intuitionist logics, are at home.

¹³ The fourth axiom, omitted here, is $\text{Con}(\emptyset) = \emptyset$.

But it is at least not obvious that structural constraints such as monotonicity and transitivity that are appropriate for logical consequence relations should be taken to hold for consequence relations more generally, including material consequence relations that do not involve logical vocabulary at all. Before looking at some examples, it will be helpful to layout a few conceptual tools that can help frame the issue. We can start naively. Suppose we have a set of sentences Γ and ask about their “content.” Their “content” is what they “contain.” (Traditional logicians, from the Port Royale logicians through Leibniz and Kant found it natural to talk this way.) Thinking of Γ as a set brings with it one obvious sense of “contain.” The set contains the sentences that are the elements of the set.¹⁴ But suppose $\Gamma \sim A$. Then there is another obvious sense in which Γ “contains” A . It *implies* A . So A can be “drawn out” of Γ , or “found in” Γ . It takes more “effort” to extract (literally “pull out”) Γ ’s consequences than it does to inspect its elements. But the metaphor of what is “in” Γ straightforwardly works for what Γ implies, too. We can call the first sense the *explicit* content of Γ . For the elements are what make it the set that it is—before we consider any reason relations that set might stand in to anything else. But if the set does stand in implication relations to other sentences, if it has consequences, if, as we have assumed, $\Gamma \sim A$, then we can say that A is *implicit* in Γ , in the literal sense of being *implied* by it. A is an *implication* of Γ and in *that* sense is “implicit” in it. So A is part of the *implicit* content of Γ , and any $G \in \Gamma$ is part of its *explicit* content. This is a specifically *semantic* conception of the explicit/implicit distinction, since it applies specifically to the *content* of sets of sentences. Note how nicely it fits with the principle of Pragmatically Implicit Acceptance, discussed at the end of Section III, according to which if commitment to accept all of Γ precludes entitlement to reject A , then it *implicitly commits* one to accept A , in a *pragmatic* sense of “implicitly.”

Understanding Γ , grasping its content, requires going beyond its explicit content, to explore also its implicit content. Figuring out what Γ implies is a kind of reasoning. But extracting consequences from Γ , in the sense of realizing what it implies, is not the same thing as *drawing conclusions* from it, if by that one means accepting A because one accepts Γ and understands that $\Gamma \sim A$. For—as Harman reminded us in enforcing the distinction between practices of *reasoning* so as to alter our commitments (and so entitlements) and reason *relations* of implication and incompatibility—one might also be committed to accept Δ , where $\Delta \# A$. And those reasons *against* A might be weightier than the reasons Γ provides *for* A . Then one should reconsider one’s entitlement and commitment to some of the elements in Γ . Extracting its implicit content can alter one’s attitude towards aspects of its explicit content.

So how should we think about reasoning in the sense of articulating the implicit content of a set of claimables by exploring implications in which it figures as the premise-set, if we are

¹⁴ This line of thought would not look too different if we considered multisets, or lists, or partially ordered sets instead of just sets.

careful to heed Harman’s warning and not to confuse that activity with reasoning in the sense of altering one’s commitments (and so entitlements) to accept and reject? Asking this practical question is one way of thinking about what it means to say that Γ implies A, what the snake turnstile means in the metavocabulary in which we write “ $\Gamma|\sim A$.” As a model we could think of the premise-set Γ as a *database*. Its explicit content, the sentences that are elements of the set Γ , are the *records* in the database. We can then think of using that database as fodder for an *inference engine*. When the inference engine (thought of as a function or program) is fed the database Γ as its argument, it computes a consequence set. (Analogous, in this non- or prelogical case, to Tarski’s “Con(Γ)—in my version.) “ $\Gamma|\sim_{IE}A$ ” then says that A is in the consequence-set of the premise-set Γ , when computed by the inference engine $|\sim_{IE}$. Since we are interested in material consequences, we are treating the sentences in the database as *logically* atomic. But that is entirely compatible with treating them as labels for records having as much internal structure as we like (perhaps specified in some object-oriented programming language). The results of applying the inference engine being considered to a particular database might crucially depend on the details of the internal structure of the records in it. The idea is that the inference engine lets users extract information from the databases it is applied to that is not *explicitly* contained in the database—that is not an entry in any record in it—but that is *implied* by what *is* there. If the database contained the record “Block #23 has the shape of a dodecahedron,” the inference engine might take that as implicitly licensing answering “Yes” to the question “Do any of the blocks have the shape of a Platonic solid?”.

Consider the Tarskian structural principle of Containment (CO) from this point of view. It says that all the elements of the premise-set follow from that premise-set. If the consequences of the premise set, are thought of as the answers to questions that it enables by means of an inference engine, surely it is entirely in order to count actual records in the database, elements of the premise set as among those consequences. After all, in the example just offered, we would want the inference engine also to answer “Yes” if asked “Does block #23 have the shape of a dodecahedron?”. And it is reasonable to take what a premise set Γ explicitly includes as also part of what it implicitly contains. This dimension of structure is one that consequence relations in general, non-logical consequence relations, seem to share with specifically logical ones.

It is quite otherwise with monotonicity (MO). Monotonicity says that if an implication is good, that it remains good under the addition of arbitrary additional premises. If $\Gamma|\sim A$, then $\Gamma,\Delta|\sim A$. There are some good candidates for general (nonlogical) implications for which this principle holds, for instance. (I express them using conditionals, but the point is the implications, which don’t depend on using logical vocabulary to codify them):

- 1) If Pedro is a donkey, then ($|\sim$) Pedro is a mammal.
- 2) If the indicated plane figure is a square, then it is a rectangle.

But there are many others that do not seem to be monotonic.

- 3) If it is raining, then ($|\sim$) the sidewalk will be wet.

But not if there are sheltering trees or awnings.

4) If Harriet will be there, then it will be a good dinner party.

But not if Sam comes, because he and Harriet are always publicly unpleasant to one another and spoil the occasion.

Most reasoning outside mathematics and fundamental physics is *probative* rather than *dispositive*. Good reasons can often be given for conclusions, but such reasons for a conclusion can most often also be contested by reasons against it. Very often, good reasons for drawing a conclusion can be defeated by providing further information. A set of reasons typically does not settle the issue of whether one should draw a particular conclusion once and for all. That the defendant pulled the trigger that fired the fatal bullet provides good reason to think that she is guilty of murder. But not if she is insane, did not believe the gun was loaded, was unaware of the presence of the victim, and so on. In the institutionalized context of forensic reasoning, the defeasibility in principle of almost every particular consideration is both what requires trials and what makes their outcomes difficult to predict or even to assess as to the correctness of their verdicts. To shift to another institutionalized reasoning context, complex medical diagnosis is an adventure in plausible, probative reasoning for conclusions that are routinely then revised or discarded upon the receipt of new, discordant information. Like the legal case, this process is the core of an entire dramatic genre, in which the audience is invited to follow the twists and turns of reasoning as new evidence and further developments require rethinking what follows from original facts that might remain fixed throughout, as the recontexting of those facts opens up and closes off various implicational pathways.

Even outside of such institutionalized reasoning practices, we can identify “Sobel sequences” of implications whose goodness is reversed *ad libitum* by the addition of further premises:

- 5) If I were to strike this dry, well-made match, then (\sim) it would light.
- 6) If I were to strike this dry, well-made match in a room with a strong magnetic field, then it would not light (# it would light).
- 7) If I were to strike this dry, well-made match in a room with a strong magnetic field and I were in a Faraday cage, then (\sim) it would light.
- 8) If I were to strike this dry, well-made match in a room with a strong magnetic field and I were in a Faraday cage, and there was no oxygen in the room, then it would not light (# it would light).
- 9) ...

One manifestation of the nonmonotonicity of general, material consequence relations is of particular significance for an important class of inference engines: those that reason probabilistically. For receipt of further information can affect the reference class with respect to which frequencies are determined in order to assign probabilities. Alvin is standing in front of what looks to him like a barn. Under ordinary circumstances, he knows what barns look like, and

so his inclination to report the visible presence of a barn makes it very likely that it is, in fact, a barn. Upon further investigation, we discover that there are hobbyists who construct *trompe l'oeil* barn façades so cunning that no-one with Alvin's standard barn-discriminating abilities can tell them from the real thing, but that in the Pennsylvania, where Alvin is, only 1 out of 100 apparent barns is a façade. The conclusion still follows from our total information that he is very likely looking at a real barn. If we acquire the further, new information that Alvin is not just in Pennsylvania, but is also in Barn-Façade County, hotbed of the façade-making hobby, where only 1 out of 100 apparent barns is a real barn, this shifting of the reference class for assessing probabilities would infirm the conclusion, making it very *unlikely* that he is looking at a real barn. However, if we then find out that Alvin is in No-Nonsense Township, which has forbidden façade-making, the conclusion would once again be instated. Until and unless we discover in addition that the farmer whose land he is on in No-Nonsense Township is a scofflaw who has been cited repeatedly for illegal barn-façade construction, which would reverse the conclusion. Nonetheless, we might still find out that what Alvin is looking at is an actual barn, built as a goodwill gesture to his neighbors by the reformed, now apostate former façade enthusiast. In short, probabilistic reasoning provides systematic resources for the construction of nonmonotonic Sobel sequences, where the goodness of the implication $\Gamma \sim A$ can be flipped by the sequential addition of further premises. The large and important class of Bayesian inference engines should accordingly be nonmonotonicity tolerant.

Unlike specifically *logical* consequence relations, as construed by Tarski and Gentzen, then, implication relations *in general* should not be presumed to be monotonic. Even logical implications can be infirmed by subtracting premises, taking away information. Material implications can also sometimes be infirmed by *adding* premises or information. And the same thing holds for incompatibility relations in general. A set of commitments that is incoherent in the sense that one cannot be entitled to all of its explicit members can be rendered coherent not just by rejecting some of them, but also sometimes by adding further commitments. Having been convicted of a felony is incompatible with (precludes legal entitlement to) later serving as a law enforcement officer—unless in addition one has been officially pardoned. Having an episode of malignant hyperthermia precludes later anesthesia with succinylcholine—unless the patient is first given dantrolene.

One consequence of the nonmonotonicity in general of nonlogical reason relations of implication and incompatibility is that the extent to which an interlocutor understands a claimable content depends not only on being able in practice to tell that certain implication or incompatibility relation involving it is a good or bad one, but also on having at least some grasp on the range of additional circumstances in which it would *remain* good or bad. Kaplan's implicational phase-space semantics captures these ranges of subjunctive robustness by assigning candidate implications \vee -sets as their semantic interpretants. These specify what additional premises would and what would not infirm a good implication, or make one that did not hold

into a good one. The limiting, degenerate case is where one endorses an implication $\Gamma \sim A$ as what Abelard calls an inference “merely *hic et nunc*,” that is, with *no* commitment to there being *anything* else that one could find out that would *not* infirm the implication. It is not clear what practical use such implications (or incompatibilities) could have, what difference they could make.

All this is of course not to deny that some reason relations *are* monotonic. Indeed, some nonlogical implications are good in virtue of their implicational structure. The Tarskian principle of Containment, which I have endorsed, says that implications of the form $A \sim A$ (instances of Reflexivity) are not just good, but are *monotonically* good. They remain good upon arbitrary additions to their premises (and, in multisuccedent settings, to their conclusions, disjunctively considered, as well). For just this reason, relevance logicians reject CO. Another subclass of implications whose monotonicity is widely endorsed because of their special structure is instances of the structural principle of Cautious Monotonicity. The idea here is that while it might indeed not be the case that some good implication $\Gamma \sim A$ remains good upon *arbitrary* additions of premises to Γ , surely it should be subjunctively robust under the addition, as auxiliary premises, of *other* consequences of Γ : claimables that it already implies.

Cautious Monotonicity (CM):
$$\frac{\Gamma \sim A \quad \Gamma \sim B}{\Gamma, B \sim A}.$$

So for example, (4) above, Harriet will be there \sim It will be a good dinner party can hold even though

10)* Harriet will be there and Sam will be there \sim It will be a good dinner party does not hold. But if

11) Harriet will be there \sim Vivian will be there,
then if (4) is good, so should be

12) Harriet will be there and Vivian will be there \sim It will be a good dinner party.

And there will be an analogous version of CM for incompatibility. Although some incompatibilities $\Gamma \# A$ might be “curable” (defeasible) by the addition of further premises, it seems plausible that one cannot alter what Γ is incompatible with by adding to it further premises that it already implies. For they are already implicitly part of the content of Γ , which is what rules out the incompatible content A .

Plausible as CM might be on its own, I think there are decisive reasons to reject it as a structural principle holding for implication (or incompatibility) in general. To see why, we can think about its complement, Cumulative Transitivity, Gentzen’s “Cut.”

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

If a set of premises Γ , together with something that it implies, B , imply A , then Γ already implies A by itself. For what we need to “add” to Γ to get A is something it already implicitly contains, part of its implicit content, Γ ’s implicational consequence B . Transitivity says that one can

safely employ the consequences of one's commitments as premises for further implications. This is the feature of mathematical reasoning that makes it possible to combine implicational links into the sort of long deductive chains that so impressed Greek philosophers and their heirs about Euclid's geometry. The possibility of using the conclusions of some pieces of reasoning as the premises from which to draw further conclusions has seemed to many to be essential to the idea of reasoning as such. Many who are willing to acknowledge the nonmonotonicity of implication draw the line at relaxing the requirement of transitivity. (This is true even of Jean-Yves Girard, the father of linear logic. He is comfortable enough with substructurality not only to reject monotonicity, but more radically, to reject Gentzen's Contraction: the principle that if $\Gamma, A, A \vdash B$, then $\Gamma, A \vdash B$. Girard takes Cut to be what makes communication possible.) Almost the only logicians willing to relinquish transitivity are those who find themselves driven to do so in order to deal with semantic paradoxes.

But let us think about CT and CM from the point of view of the two bits of conceptual apparatus I have suggested as helpful for addressing questions of the structure of reason relations in general. These are using implication relations to distinguish between the explicit and implicit contents of a set of premises—what they contain explicitly and what they contain implicitly—and the database-plus-inference-engine model of implication. The first lets us see the sense in which CT and CM are complementary principles. For they both concern the effects of what we can call “explicitation.” This is taking part of the *implicit* content of a set of premises and adding it to the *explicit* content: making that implicit content explicit. Since the implicit content of a premise set is what is implied by it and its explicit content is the elements of that premise set, this is moving implied consequences from the right-hand side of the implication turnstile to the left-hand side. In these terms, CT says that explicitation does not *add* any implicit content, and CM says that explicitation does not *subtract* any implicit content.

For CT says that if B is implicit in (implied by) Γ , and the result of adding it to Γ explicitly, Γ, B (shorthand for $\Gamma \cup \{B\}$) implies A, then Γ already implied A. Explicitating the implicit content B did not result in any new consequences, any new implicit content. And CM says that if B is implicit in (implied by) Γ , and Γ implies A, then so does the result of adding B explicitly to Γ as an additional premise. The implication of A by Γ is not defeated by making explicit the additional bit of Γ 's implicit content that is B. So together, CT and CM say that explicitation, moving a sentence from the right to the left of the implication turnstile, does not affect the implicit content of a set of premises at all. Changing the form of the content from implicit to explicit makes no difference to the implicit content. That is, CT and CM together entail the *inconsequentiality of explicitation*. Making explicit the implicit content of a premise set does not change the implication potential of that premise set.

Now we can ask whether it makes sense from the point of view of the database-plus-inference-engine model of implication to impose the inconsequentiality of explicitation as a global structural principle on implications in general. In the database setting, the

inconsequentiality of explicitation entailed by the conjunction of CT and CM says that one *must* treat the consequences extracted from the database by applying an inference engine to it as having *exactly the same status* as the records explicitly contained in the database. One can add those implications to the database confident that doing so will make *no difference* to the implicit content and consequences of that database. But is this a good idea? Suppose the database contains the observations made so far using CERN's Large Hadron supercollider, and the inference engine applied to that data to extract its consequences is the Standard Theory of contemporary particle physics. There are many predictions that the Standard Theory makes from current observations that it would be a big deal to have corroborated by observation. The principle of the inconsequentiality of explicitation says on the contrary that empirical confirmation of theoretical predictions from observation is trivial in principle, that it can never make an epistemic difference. The inconsequentiality of explicitation would not only permit entering the results of theoretical predictions (the results of applying the inference engine of the Standard Theory to the observations in the database) into the database with the same status as actual observations, but would require that doing so make no difference. It elides the distinction between empirical observations and theoretical predictions that is what makes room for the very idea of empirical confirmation of theories.

Thinking about the issue from this point of view makes clear that imposing CT and CM as structural principles governing reason relations generally prohibits reasoning according to any sort of implication that allows the possibility of according a different epistemic status to the implicational consequences of a set of premises than are accorded to those premises themselves. For it is only when there is some such difference in epistemic status that explicitation will *not* be inconsequential. Any sort of risky inference, the inferential employment of any theory whose correctness is less certain and secure than the results of observation, would be forbidden. That is incompatible with essentially every institutionalized reasoning practice outside of mathematics and logic. In science, we have seen, it would make the notion of empirical confirmation of theories unintelligible, and this is no less true in its applications in engineering and medicine. In law, eliding the distinction between evidence and the results of applying legal reasoning to that evidence amounts to transcendentally deducing that appeals courts are superfluous or nugatory. We have to be able to treat the results of some of our inferences as less solid than the evidence from which we made those inferences, and so be conscientious about the grounds we have for promoting theoretical conjecture to the status of established fact.¹⁵

It follows from these considerations that reason relations in general, material, nonlogical relations of implication and incompatibility, exhibit little of the structure characteristic of specifically *logical* implication and incompatibility (inconsistency). Of the Tarskian structural principles defining logical consequence as a topological closure operator, with which we began,

¹⁵ This is, of course, the leading idea of the strict-tolerant semantics for logical vocabulary, resulting in the logic ST mentioned in the previous section. [ref. to Barrio et. al.]

only Containment remains as plausibly characterizing material consequence relations in general. From the traditional logical point of view, the reason relations governing everyday discursive practice (as well as their institutionally regimented descendants in specialized domains such as empirical science, medicine, and law) are radically substructural.¹⁶

¹⁶ I offer a different argumentative path through the material of this and the next sections, in “From Logical Expressivism to Expressivist Logic: Sketch of a Program and Some Implementations” *Nous: Philosophical Issues*, Volume 28, Issue 1, October 2018, (a volume devoted to the philosophy of logic), pp. 70-88.

VI. Logic and the Explicit Expression of Substructural Reason Relations

Practices of giving and asking for reasons are practices of making, defending, and challenging claims. Claims are bipolar commitments that count as specifically *doxastic* commitments because of the way their entitlements are always potentially at issue, able to be challenged, and needing to be redeemed or vindicated by reasons. Those commitments and entitlements are articulated by two sorts of opposite but intimately related reason relations. Implications are relations of being a reason for, appealed to in rationally defending basic doxastic commitments, and incompatibilities are relations of being a reason against, appealed to in rationally challenging those commitments. The different roles that implications and incompatibilities play in reasoning practices requires that the two kinds of reason relations have different structural properties. Incompatibilities are *de jure* symmetric, while implications in general are not symmetric. Of the topological structural properties that Tarski argued characterize *logical* implication, implication in the general, nonlogical case satisfies only the principle that the explicit premises of an implication are included in the consequences that they imply. Their implicit content, what they imply, includes their explicit content, the premises that imply those consequences. Nonlogical implications are not in general monotonic, or even cautiously monotonic, nor are they in general transitive—though special cases sometimes do conform to these structural principles. Like implications, material incompatibilities are not in general monotonic. Sometimes, adding further auxiliary hypotheses to a premise set can make something that was incompatible with it compatible instead (as well as the other way around).

In articulating these relations and making these arguments, I have not found it necessary to say anything about specifically *logical* reason relations, except to contrast their rigid structure with the more relaxed structural conditions that apply to reason relations outside the realm of logic and mathematics. But there is a once dominant and still influential tradition that says this way of approaching things must be wrong. Once upon a time there was an ideology I will call “logicism about reasons.” According to this view, all genuine reasons are at base logical reasons, in the sense that what makes a reason a *good* reason is that there is a *logically valid* argument connecting premises and conclusion.¹⁷ Underneath or behind every material incompatibility there is a logical inconsistency. (Compare the kind of view that Dewey contrasts with his pragmatism as Platonism or intellectualism, which finds a rule or principle behind every propriety of practice.)

¹⁷ Logicism about reason relations only addressed doxastic or theoretical reason relations. It was silent on the topic of practical reasons, and so had nothing to say about their credentials to be called “reasons” as falling under a genus that has both doxastic and practical species.

According to this view, Aristotelian contraries are to be understood in terms of conceptually prior Aristotelian contradictories, which essentially involve the *logical* concept of negation. So “T is square” is incompatible with “T is circular” in the sense that and because “T is square” implies the *negation* of “T is circular.” To understand “material” incompatibilities, one must at least implicitly understand negation and logical contradictoriness: the way a claimable *logically* excludes its negation. Just so, an implication such as

13) It is raining. So: the sidewalk will be wet.

should be understood as good, its conclusion as following from its premise, only if understood as as enthymematic: as suppressing a missing premise. To show its goodness, one needs to supply the missing premise, and exhibit it as an instance of a logically valid schema such as *modus ponens*:

14) If it is raining, then the sidewalk will be wet. It is raining. So: the sidewalk will be wet.

To understand such implications, one must at least implicitly understand conditionals and detachment from them. Thinking along these same lines already led Sextus Empiricus to attribute mastery of the disjunctive syllogism to the dog (described by Chrysippus) who, losing sight of a rabbit he is chasing and coming to a fork in the trail, sniffs in vain at one path and then immediately runs down the other without further investigation. Surely, the thought is, the dog must have reasoned: “The rabbit either took fork A or fork B. It did not take fork A. So it must have taken fork B.”

Logicism about reasons is the ideology that motivates the idea that philosophical analysis can dig down below the surface of our unobtrusive natural languages and recover the hidden logical structure in virtue of which our good reasons are good reasons. (The *Tractatus* has scriptural status for this movement—at least as its New Testament.) But the substructural character of reason relations in general—when we do not restrict them to *logical* implications and incompatibilities-as-inconsistencies—shows that this cannot be right. The program of turning ordinary defeasible implications and incompatibilities into, or deriving them from, nonmonotonic logical ones is unworkable. We use defeasible reason relations because the reasons we can actually give must be finitely storable, and the number of potential defeasors is unlimited. They cannot all be ruled out explicitly by adding additional premises. Worse, the class of potential defeasors is not just *infinite*, which might be dealt with by judiciously chosen generalizations, but *indefinite*. There is no way to survey all the possible conditions that could defeat an implication such as “This is a dry, well-made match. So it will light if struck.” We might think it possible in principle to completely survey things that could go wrong from the point of view of physics and engineering (gravitational fluctuations, being doused by a bucket of water...), but what about the match turning into well-cooked pasta, or interfered with by spells cast by elves, arbitrary changes in the laws of nature...? The attempt explicitly and exhaustively to specify the range of subjunctive robustness of an ordinary implication like this is hopeless. And the open-ended nature of possible social complications surely illustrates the point even more

clearly for “Harriet will be there, so the dinner party will be a success,” or “The Fed is increasing the money supply, so the stock market will rise.”

Some have thought to employ *ceteris paribus* clauses to do at a stroke what admittedly cannot be spelled out in detail. But this is totally to misunderstand the expressive function of such clauses. The use of a *ceteris paribus* “all things being equal” clause just explicitly *acknowledges* the nonmonotonicity of the implication, it does not *remove* it. (The term for a Latin phrase whose utterance could do *that* is “magic spell”—and we should leave them to the mischievous possible elves.) Only on the defeasibility-acknowledging reading of *ceteris paribus* clauses is their use legitimate. Otherwise saying “*q* follows from *p*, *ceteris paribus*” would be saying “*q* follows from *p*, except in cases where for some reason, it doesn’t.” And that is empty.

So logical reason relations exhibiting the full Tarski-Gentzen structure cannot underlie material reason relations in general, in the sense of being what distinguishes between good reasons and bad reasons. Is the proper conclusion, then, that traditional logics are, in virtue of their structural commitments simply irrelevant to reason relations in general, because of the radically substructural character of material relations of implication and incompatibility? Is it the case that traditional logic applies at most to the reflexive, monotonic, transitive reason relations of mathematics? That view is not without its advocates. But it is predicated on a mistaken view of the relations between logical and material reasons. Logic is not what *makes* nonlogical reasoning good (when it is good). The proper task of logical vocabulary is rather to let us *say* what reasons are good, what follows from what and what is incompatible with what: to make it possible for us to *talk about* what follows from and rules out what else. It is to bring those reason relations into the discourse, in the form of claimables that can themselves be accepted and rejected, defended and challenged, as things reasons can be given for and against.

Instead of understanding the significance of logic by the project of looking for something displaying the “hardness of the logical ‘must’” underneath the norms governing our ordinary reasoning, we should look to the distinctive *expressive* role characteristic of logical vocabulary. Exploring and exploiting the reason relations that bind claimables together and articulate their content is how we assess the reasonableness of our doxastic attitudes. Logical vocabulary lets us codify those reason relations themselves into claimable form, and so to be critical about them. It is what makes it possible to have reasons to change what we count as reasons. And since material reason relations of implication and incompatibility articulate the claimable contents that can be doxastically accepted or rejected, to which interlocutors can be committed and, if all goes well, entitled, using logical vocabulary that plays the expressive role of making them explicit as claims that can themselves be challenged and defended makes possible a distinctive kind of rational semantic self-consciousness.

Demarcating distinctively *logical* vocabulary (and the class of distinctively logical concepts it expresses) by the expressive role it plays in making reason relations explicit means that the essential, paradigmatic logical locutions are the conditional and negation. For conditionals serve to make implications explicit and negation serves to make incompatibilities explicit. Accepting “ $A \rightarrow B$ ” is committing oneself to the goodness of the implication of B by A: what we have been expressing in the metavocabulary of reason relations as “ $A | \sim B$.” Accepting “ $\sim A$ ” is committing oneself to something incompatible with A, something that provides a reason against A, a reason to reject A. More specifically, to perform its expressive task of codifying incompatibility relations in the object language, negation needs to satisfy the

Minimal Negation Condition (MNC): $\Gamma | \sim \sim A$ iff $\Gamma \# A$.

That is, a premise-set implies not-A just in case A is incompatible with that premise-set.

Correspondingly, to perform its defining expressive task of codifying implication relations in the object language, conditionals need to satisfy the

Ramsey Condition (RC): $\Gamma | \sim A \rightarrow B$ iff $\Gamma, A | \sim B$.

That is, a premise-set implies a conditional just in case the result of adding the antecedent to that premise-set implies the consequent. A conditional that satisfies this equivalence can be called a “Ramsey-test conditional,” since Frank Ramsey first proposed thinking of conditionals this way.

The Minimal Negation Condition says that $\sim A$ is the *least incompatible* of A, in the sense of being what is implied by anything that is incompatible with A. Where logicism about reasons wants to define Aristotelian material contraries by Aristotelian logical contradictories, that order of explanation is reversed here. Red’s contradictory, not-red is defined as what is implied by every contrary of red: so, by blue, yellow, and so on. Similarly, the Ramsey Condition says that $A \rightarrow B$ is what is implied by everything that, when combined with the antecedent A of the conditional, implies the consequent B of the conditional. The conditional is defined as implicationally the weakest sentence that has this property. In any particular case there might already be an atomic sentence implied by everything incompatible with some particular sentence, or for some pair there might be something implied by everything that, together with the first, implied the second. But introducing the connectives gives a systematic method for producing sentences guaranteed to play these roles, even if the language is extended by the addition of new atoms. (Semantically, these are both instances of defining connectives by residuation relative to incompatibility and implication relations, respectively.)¹⁸

By specifying necessary and sufficient conditions on the roles in implications that locutions must play in order to count as conditionals or negations, the Ramsey Condition and the Minimal Negation Condition go some way to filling in the idea that conditionals *express*

¹⁸ This expressivist account takes a frankly inegalitarian attitude toward different classes of sentential logical connectives. In particular, conjunction and disjunction play a secondary expressive role, as mere combinatory helper-monkeys codifying what is expressed in the metavocabulary of multisuccedent sequent calculi by commas on the left and the right side of turnstiles, respectively.

implication relations and negations *express* incompatibilities. For those conditions specify that one is implicitly committed to conditional or negated claims just in case one endorses certain reason relations involving their nonlogical components. They do that by saying what reason relations one needs to take a premise set Γ to stand in to other sentences in order for it to imply conditionals or negations formed from them. The RC and MNC provide cash for the claim that conditionals make it possible for one explicitly to *say that* an implication or incompatibility holds. This permits an important kind of *semantic descent*: reason relations that must otherwise be expressed in a metavocabulary (such as the one we have been using, with ‘ \sim ’ expressing implication and ‘ $\#$ ’ expressing incompatibility) can now be expressed in the logically extended object language. The reason relations that articulate the claimable (propositional) contents of the object language can now be expressed themselves as claimable contents, for which reasons can be given and demanded.

The picture here is that we start with a language L_0 that has no logical vocabulary: a set of logical atoms. (Whatever other lexical-syntactic articulation or record-structure its sentences might have, they are not *logically* complex.) We construe the sentences of that language as standing to one another in material reason relations of implication and incompatibility satisfying only the structural constraints that implication satisfies CO, and incompatibility is symmetric.¹⁹ The pair of an implication relation and an incompatibility relation for the language can be called a “material semantic frame” (MSF) defined on it. The base language L_0 can be extended lexically and syntactically into the logically extended language $L_{\rightarrow\sim\&\vee}$ (or just L , for short), which is the smallest set that contains L_0 and for any two sentences A and B of L contains $A\rightarrow B$, $\neg A$ and $\neg B$, $A\&B$, and $A\vee B$. The dual semantic criteria of adequacy for a set of rules determining the implication and incompatibility relations governing the logically complex sentences of L is that for any MSF_0 defined on L_0 they determine a corresponding MSF_L on the logically extended language L that is i) a conservative extension of the original MSF_0 and ii) permits the expression in L of all the implications and incompatibilities that hold in that extended MSF_L .

Gentzen’s sequent calculi show how to formulate connective rules that do just that for all prior MSF_0 s that are “flat” in the sense that all of their implications are instances of Reflexivity, that is, are of the form $A|\sim A$. (The result is classical logic if the implications are construed as having multiple conclusions, and intuitionist logic if the implications are construed as restricted to single conclusions.) And those same connective rules can be applied to substantive sets of nonlogical “axioms”—as Gentzen himself does when using his sequent calculi to investigate the consistency of Peano arithmetic. So from one point of view, the expressivist approach to logic just redescribes the founding achievements of the proof-theoretic tradition. A substantial difference becomes visible, however, when we ask about the background assumptions about the structure of the reason relations that show up when Gentzen’s substantive (non-RE) “axioms” play the role of the base

¹⁹ We’ll actually also want to require that the set of all the sentences in the language be incoherent, so any partition of partitions it into incompatible sets. But I won’t make any use of that condition here.

MSF₀. For Gentzen requires that the axioms be fully structural, and his rules preserve that structure: a monotonic, transitive base will yield a monotonic, transitive logical extension. Relax those structural conditions at all and the system blows up. (The proof of Gentzen’s “Hauptsatz” is the first casualty, but it that is just the beginning of the troubles. Gentzen’s connective definitions—for instance, his rules for conjunction—force monotonicity, and so will fail even to be conservative when applied to a nonmonotonic base MSF₀.) Gentzen’s connective rules cannot be used to codify substructural material reason relations. That is, they won’t work for any that are not uniformly monotonic and transitive, and so build in the inconsequentiality of explicitation.

Remarkably, however, Hlobil and Kaplan have shown that connective definitions that are in a clear sense *equivalent* to Gentzen’s *can* satisfy the dual semantic criteria of adequacy for logical vocabulary even in radically substructural settings. How can that be? The variants of Gentzen’s connective rules that they identify are equivalent to his *in his fully structural setting*. So, Gentzen himself could have done everything he did using these variant formulations. Indeed, they are essentially the variants that his student Ketonen formulated in his dissertation, to make Gentzen’s connective rules reversible (to make them into “double horizontal line” rules). But sets of connective definitions that are equivalent in strongly structural situations can come apart in substructural settings, defining quite different extensions of the same underlying material MSF₀s. That is what happens here. Hlobil and Kaplan’s structurally equivalent but substructurally divergent Ketonen-variants of Gentzen’s rules are fully conservative over substructural material base implications-and-incompatibilities, and still, like Gentzen’s own, permit the explicit expression, by sets of sentences in the logically extended language, of all the implication and incompatibility relations, not only of the base MSF₀, but also of its extension to the full logically extended language.

Here's how to define the system they call “NM-MS” (for NonMonotonic Multi-Succedent):

1. Begin with a language L_0 consisting of a finite number of logically atomic sentences, and a basic material semantic frame MSF_0 consisting of an implication relation $|\sim_0 \subseteq \mathcal{P}(L_0) \times \mathcal{P}(L_0)$. We encode material incompatibility into $|\sim_0$ by stipulating that if $\Gamma|\sim$, that is, the right-hand side of the snake turnstile is empty, then Γ is incoherent, and any two sets whose union is Γ are incompatible with one another.
2. Extend L_0 to include logically complex sentences formed from L_0 , defining L as the smallest set by inclusion that contains L_0 , and for any two sentences $A, B \in L$ contains also $A \rightarrow B$, $\neg A$, $\neg B$, $A \& B$, and $A \vee B$.
3. Extend MSF_0 on L_0 to a material semantic frame MSF on L by extending $|\sim_0$ to $|\sim$.

Stipulate first:

Axiom: If $\Gamma_0|\sim_0 \Delta_0$ then $\Gamma_0|\sim \Delta_0$.

Then close $|\sim$ under the following rules:

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

$$\text{R}_{\rightarrow}: \frac{\Gamma, A \mid \sim \Delta, B}{\Gamma \mid \sim \Delta, A \rightarrow B.}$$

$$\text{L}_{\neg}: \frac{\Gamma \mid \sim \Delta, A}{\Gamma, \neg A \mid \sim \Delta.}$$

$$\text{R}_{\neg}: \frac{\Gamma, A \mid \sim \Delta}{\Gamma \mid \sim \Delta, \neg A}$$

$$\text{L}_{\&}: \frac{\Gamma, A, B \mid \sim \Delta}{\Gamma, A \& B \mid \sim \Delta.}$$

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

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

$$\text{R}_{\vee}: \frac{\Gamma \mid \sim \Delta, A, B}{\Gamma \mid \sim \Delta, A \vee B.}$$

This system is very similar to the system Negri, Von Plato, and Ranta call “G3cp.”²⁰ The difference is that NM-MS uses Gentzen’s (and Ketonen’s) more familiar rules for negation in place of G3cp’s. Ketonen’s innovation was to formulate rules that are reversible, in the sense that the moves from what is below the horizontal line to what is above it are all provably admissible. So, for instance, any context that permits the derivation of $\Gamma \mid \sim \Delta, A \rightarrow B$ also permits the derivation of $\Gamma, A \mid \sim \Delta, B$ (the converse of R_{\rightarrow}), and any context that permits the derivation of $\Gamma \mid \sim \Delta, \neg A$ also permits the derivation of $\Gamma, A \mid \sim \Delta$. This reversibility is important because it means that the *biconditional* Ramsey Condition on conditionals and the *biconditional* Minimum Negation Condition both hold in NM-MS. Satisfying these conditions, I argued earlier, is essential to sentences formed using these connective performing their defining expressive function of codifying implications and incompatibilities, respectively.

²⁰ Negri, S., Von Plato, J., & Ranta, A. *Structural Proof Theory* [Cambridge University Press, 2008]. The difference is that NM-MS uses Gentzen’s (and Ketonen’s) rules for negation. The sets of connective rules are equivalent in the sense that G3cp’s negation rules are admissible in NM-MS, and NM-MS’s negation rules are admissible in G3cp. In a fully structural setting, G3cp is a sound and complete proof system for classical logic.

These connective rules conservatively extend even substructural MSF_0 s, defined on L_0 to correspondingly substructural MSFs defined on the language L of logically complex sentences formed from those atoms. For this to be true, it is essential that $L_{\&}$ and R_{\vee} take the “multiplicative” form that they do. For if the left rule for conjunction, for instance, were the “additive”

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

its reversibility would not in a nonmonotonic MSF support the inference

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

since in such a setting that conclusion does not in general follow from $\Gamma, A \mid \sim \Delta$ and $\Gamma, B \mid \sim \Delta$.

The connective rules of NM-MS ensure conservativeness over material semantic base relations of implication and incompatibility that are not everywhere monotonic and transitive. By satisfying RC and MNC, they permit the explicit expression of implications $\Gamma \mid \sim \Delta$ in the logically extended object language by sentences of the form $(G_1 \& \dots \& G_n) \rightarrow (D_1 \vee \dots \vee D_m)$ (where the G_i are the elements of Γ and the D_j are the elements of Δ) and of incompatibilities of the form $\Gamma \# A$ by sentences of the form $(G_1 \& \dots \& G_n) \rightarrow \neg A$. But the rules of NM-MS make it possible to express in the logically extended object language the prelogical reason relations that hold in the base MSF_0 to which they are applied in a much more striking and complete manner than this. For Kaplan has proven a remarkable representation theorem relating theories (sets of sentences) in the logically extended language to arbitrary sets of implications and incompatibilities in the material semantic frame defined on the atomic base language, to which the NM-MS rules are applied.²¹ For any arbitrary set of sentences in L , it is possible to compute exactly what consequences have to hold in the base MSF_0 in order for that logically complex theory to be derivable from the base. And for any arbitrary set of consequences in the base MSF_0 , it is possible to compute what set of sentences of L will be derivable in any MSF_0 in which they hold. When the logical vocabulary is defined as it is by NM-MS, it makes it possible to *say* in the logically extended object language what implications and incompatibilities hold in the base language—and for that matter, in the logically extended language defined on that base. With the aid of the combining connectives $\&$ and \vee , NM-MS makes it possible to use conditionals and negations *say* what we have been specifying in the metalanguage of reason relations by using ‘ $\mid \sim$ ’ and ‘ $\#$ ’.

Furthermore, NM-MS is an exceptionally well-behaved logic. If it is applied to a base or implicational prior that is “flat”—in the sense of consisting, like Gentzen’s axioms, entirely of

²¹ Dan Kaplan “A Multi-Succedent Sequent Calculus for Logical Expressivists” in Pavel Arazim and Tomas Lavicka (eds.) *The Logica Yearbook 2017* [College Publications, 2018].

instance of Reflexivity, so of the form $A \mid\sim A$ —then it just yields classical logic. Further, the MSFs that result from applying the rules of NM-MS to arbitrary substructural (nonmonotonic and nontransitive) bases are guaranteed to be supraclassical—they validate all the *theorems* of classical logic—even though their extended *consequence* relations and incompatibility relations (which are being encoded as implications with empty right-hand sides) are quite different, in being nonmonotonic and nontransitive. Nonetheless, the *logical* consequence relations and incompatibility relations, those that hold no matter what underlying base MSF₀ they are defined on, that is, that hold just in virtue of the logical vocabulary they contain (in the sense of holding under arbitrary substitutions of nonlogical for nonlogical vocabulary) are *de jure* monotonic and transitive. NM-MS is a logic of nonmonotonic and nontransitive consequence and incompatibility relations, *not* itself a nonmonotonic or nontransitive *logic*. In this dual sense, NM-MS is exceptionally well-behaved even when applied to radically substructural base reason relations. Of course, the whole consequence relation (as opposed to its restriction to purely *logical* consequences—those that hold just in virtue of the connective rules of NM-MS) on the logically extended language is in general nonmonotonic and nontransitive, as it must be in order to be conservative with respect to generally nonmonotonic and nontransitive base reason relations.

So NM-MS performs the characteristic *expressive* task of logic ideally well. Using it to introduce logical vocabulary into a language-in-use that does not already contain such vocabulary permits the explicit expression, in the logically extended language, of the material reason relations of implication and incompatibility in virtue of which the expressions of the language mean what they do—play the roles in reasoning that they do. And, as a bonus, it also in the same sense makes it possible to make explicit in the same sense—to put into the form of claimables that themselves stand to one another in reason relations of implication and incompatibility—the reason relations that govern the newly introduced sentences that do contain logical vocabulary. Unlike traditional logical vocabulary, they can perform this crucial expressive function for the generally nonmonotonic, nontransitive, substructural material reason relations that actually govern reasoning with nonlogical vocabulary. Yet in fully structural settings, NM-MS is equivalent to Gentzen’s LK, that is, to classical logic. This, I want to say is what was always right and important about classical logic, the reason it deserved its pre-eminent position—*not* the fact that it can be given a particularly simple bivalent semantics in terms of the two basic truth-values. Of course, that it can is not just a bizarre coincidence, and I have been concerned here to sketch some of the more complex, ultimately pragmatic, bipolarities that semantic bivalence reflects without adequately expressing or articulating.

NM-MS was specified here proof-theoretically, by connective rules formulated in the metalanguage of Gentzen’s sequent calculi. In its general, radically substructural form, it is not amenable to a bivalent semantics. At this point it should come as no surprise that Dan Kaplan has proven soundness and completeness theorems for NM-MS with respect to his implicational

phase-space semantics for the fully general substructural case.²² (In fact, he shows how to make both the proof-theoretic and phase-space semantic characterizations work and work together in even more radically substructural settings than I have discussed here: relaxing not only monotonicity and transitivity, but CO and RE, and even Gentzen’s Contraction.) Recall from Section IV that the implicational phase-space semantics works on a space whose points are candidate multi-succedent implications, equipped with a commutative monoid of *adjunction* (a kind of fusion) defined by $\langle \Gamma, \Delta \rangle \cup \langle \Pi, \Theta \rangle =_{df.} \langle \Gamma \cup \Pi, \Delta \cup \Theta \rangle$, and a distinguished set **I** of *good* implications. Candidate implications (and sets of them) are semantically interpreted by the \vee -function, which assigns each candidate implication $\langle \Gamma, \Delta \rangle$ the set $\langle \Gamma, \Delta \rangle^\vee$ of candidate implications that, when adjoined to it, yield a *good* implication (element of **I**). The \vee -set of good implications is their range of subjunctive robustness: the premises (and conclusions) that can be added to them without infirming them. The \vee -set of other candidate implications are their good-making conditions: what one needs to adjoin to them to *make* them good. (Compare: truth conditions.) Sentences $A \in L$ are represented by pairs of candidate implications: $\langle \langle A, \emptyset \rangle, \langle \emptyset, A \rangle \rangle$. Semantically, the proposition $[[A]]$ expressed by the sentence A is

$$[[A]] =_{df.} \langle \langle \{A\}, \emptyset \rangle^\vee, \langle \emptyset, \{A\} \rangle^\vee \rangle$$

We can write the elements of this ordered pair as $[[A]]_P$ and $[[A]]_C$ (the P and C being mnemonic for premissory and conclusory roles).

In these terms, the semantic definitions that go with the connective sequent rules of NM-MS are these:

$$[[A \rightarrow B]] =_{df.} \langle [[A]]_C \cap [[B]]_P, (([[A]]_P)^\vee \cup (([[B]]_C)^\vee)^\vee \rangle,$$

$$[[\neg A]] =_{df.} \langle [[A]]_C, [[A]]_P \rangle,$$

$$[[A \& B]] =_{df.} \langle (([[A]]_P)^\vee \cup (([[B]]_P)^\vee)^\vee, [[A]]_C \cap [[B]]_C \rangle,$$

$$[[A \vee B]] =_{df.} \langle [[A]]_P \cap [[B]]_P, (([[A]]_C)^\vee \cup (([[B]]_C)^\vee)^\vee \rangle.$$

The semantic entailment relation, which parallels the horizontal line of sequent derivation, is:

$$A \models B \quad \text{iff}_{df.} \quad (([[A]]_P)^\vee \cup (([[B]]_C)^\vee)^\vee \subseteq \mathbf{I}.$$

The mapping between these settings that I provided in the Section IV shows that this is *exactly* parallel to the consequence relation Hlobil defines in the truth-maker setting, in order to parallel Restall and Ripley’s normative pragmatic bilateralism: a consequence holds just in case the fusion of verifiers of all the premises with any falsifiers of the all the conclusions is an impossible state.

The general implicational phase-space semantics is radically substructural in imposing no structure beyond what derives from the set-theoretic definition of the adjunction monoid \cup . Additional structural conditions that can be added to the NM-MS connective definitions in the sequent calculus can be mirrored exactly in the implicational phase-space (IPS) semantics, in the obvious ways:

²² What follows just transcribes Dan Kaplan’s definitions in “Substructural Content” 2019-03-05. [Better ref. ?]

Reflexivity (RE): $\forall A \in L [\langle A, A \rangle \in \mathbf{I}],$

Containment (CO): $\forall \Gamma, \Delta \subseteq L [\Gamma \cap \Delta \neq \emptyset \Rightarrow \langle \Gamma, \Delta \rangle \in \mathbf{I}].$

Monotonicity (MO): $\forall \Gamma, \Delta \subseteq L [\langle \Gamma, \Delta \rangle \in \mathbf{I} \Rightarrow \forall \Pi, \Theta \subseteq L [\langle \Gamma, \Delta \rangle \cup \langle \Pi, \Theta \rangle \in \mathbf{I}],$

Transitivity (CT): $\forall \Gamma, \Delta \subseteq L \forall A \in L [(\langle \Gamma \cup \{A\}, \Delta \rangle \in \mathbf{I} \text{ and } \langle \Gamma, \Delta \cup \{A\} \rangle \in \mathbf{I}) \Rightarrow \langle \Gamma, \Delta \rangle \in \mathbf{I}].$

Kaplan shows that imposing any combination of these structural condition on the set of good implications \mathbf{I} of a set of phase-space models yields a semantics with respect to which NM-MS is sound and complete when the same combination of structural conditions are imposed on the sequent calculus in which NM-MS extends a base material semantic frame (which must, of course, satisfy the same combination of structural conditions).

These results mean that the expressive power to codify semantogenic reason relations of implication and incompatibility that NM-MS adds to the object language by extending it (and its reason relations) by the addition of logical vocabulary is just the same as the expressive power of the implicational phase-space semantic metavocabulary to codify those reason relations. Unlike standard logics and traditional semantics, this expressive power extends perfectly smoothly to massively substructural reason relations, which I have argued are the kind we find in the wild, implicit in nonlogical, nonmathematical reasoning, both informal and in institutionalized settings such as medicine, law, and wherever inferences are drawn from databases. Recall that in Section IV I showed that any reason relations of implication and incompatibility that can be codified by the apparatus of Fine's truth-maker semantic metavocabulary can also be codified in the implicational phase-space semantic metavocabulary, and *vice versa*. I did not show there that the various structural restrictions that we have been considering can also be mirrored in the truth-maker setting, but Ulf Hlobil has shown that this is so. So there is every reason to believe that that semantic metavocabulary, too, can be turned to the expression of nonmonotonic and nontransitive material relations of consequence and incompatibility.

Part Three: The Dialogic Pragmatic Structure of Reasoning Practices

VII. “And That’s All, Folks”: Reason Relations and the Dialogic Pragmatics of Reasoning

The pragmatics-first account I have been articulating here is in many ways inspired by Restall and Ripley’s bilateralism. They have explicitly endorsed understanding the meanings of declarative sentences in terms of the role those sentences play in reason relations. To begin with, these are identified with the implication relations codified in the sequents manipulated by Gentzen-style sequent calculi. Ultimately, those implications are themselves explained in terms of the holistic property of material “incoherence” in virtue of which some “positions” consisting of constellations of affirmations and denials are “out of bounds.” Taking this step is already acknowledging one of the points I have been concerned to make here: that material, prelogical material incompatibility relations as well as prelogical material consequence or implication relations play an essential role in articulating the contents expressed by declarative sentences. Rejection of what I have called “logicism” about these basic reason relations is a core tenet of this bilateralist understanding of them. Ripley is very clear on this point: “It is from this prior ability to treat things as incoherent that negation gets its content, not vice versa.”²³

Crucially, the notion of out-of-boundness by which semantogenic reason relations are explained is a *normative, pragmatic* concept. It applies to the practical attitudes of doxastic acceptance and rejection that are expressible by speech acts of affirmation and denial. And it denominates some constellations of such attitudes as improper or inappropriate. That characterization shows itself as essentially normative in its practical bearing on conduct. One who finds himself in a doxastic position that is ruled out is both subject to negative criticism for it by others (along the second- or third-personal normative dimension of assessment) and is positively obliged to *do* something, to alter some attitudes, so as to result in a position that is back in bounds (along the first-personal normative dimension of deliberation).

This bilateralist metaconceptual analysis forms the base from which the lines of thought I have outlined here proceed. The normative pragmatic account of the reason relations that I elaborated in Part One of this essay was intended to build on and further develop the Restall-Ripley bilateralist reconstruction of them. The key was distinguishing commitments from entitlements on the side of practical normative statuses. The construction I offered (adapted from Simonelli) exploits the interaction of that distinction with the basic bilateral distinction of attitudes of acceptance and rejection to offer a more fine-grained account of implication and incompatibility relations. When I first addressed semantics, at the beginning of Part Two, the

²³ “Bilateralism, Coherence, and Warrant” in *Act-Based Conceptions of Propositional Content: Contemporary and Historical Perspectives*, Friederike Moltmann and Mark Textor (eds.) [Oxford University Press, 2017], p. 310.

point of departure was the remarkable, deeply illuminating isomorphism Ulf Hlobil constructs between Restall and Ripley's normative pragmatic bilateralist understanding of reason relations and their rendering in Kit Fine's truth-maker semantic framework. That result was then extended to Dan Kaplan's implicational phase-space semantics, based natively on implications and incompatibilities, showing how to specify an isomorphism between the two semantic frameworks concerning exactly their rendering of the reason relations Hlobil showed how to understand in Restall-Ripley bilateralist pragmatic terms.

Having demonstrated this consilience, and so vindicated the pragmatics-first approach to semantics via reason relations of implication and incompatibility—stopping along the way to explore how logical vocabulary can perform its distinctive expressive role of making those reason relations explicit in the (conservatively extended) object language, even in the massively substructural settings we find in actual discursive practice—it is time to return to consideration of the pragmatic metavocabulary in which the use of the underlying prelogical material base vocabulary is specified. For the additional structure we have introduced goes beyond what Restall and Ripley's bilateralism can accommodate. We need to be clear about the conception of discursive practice on which the rest of the conceptual edifice rests. In addition to construing implications in terms of commitment to accepting all the premises precluding entitlement *to* reject the conclusion and incompatibilities in terms of commitment to accepting all the premises precluding entitlement *to* accept the conclusion, I have talked about implications as functioning as reasons *for* claimables and incompatibilities as functioning as reasons *against* them. And those reason relations were in turn connected to things interlocutors could be understood as doing in addition to expressing their attitudes of commitment by speech acts of asserting and denying, namely *challenging* another's commitment and *defending* one's own. These speech acts change the *entitlements* interlocutors have to the commitments they undertake.

The thought is that one can *challenge* another's assertion by offering a reason *against* it in the form of endorsing premises *incompatible* with it. That then removes the default entitlement that commitment would otherwise enjoy. And one can *defend* one's own (challenged) assertion by offering a reason *for* it, in the form of endorsing premises that *imply* it. That then reinstates the default entitlement for one's assertion that was removed by the challenge. Dually, one can challenge another's denial by offering a reason for it, by asserting premises that imply it, and defend one's own denial by offering a reason against the claimable, by asserting premises that are incompatible with it. It is essential to the story I have told here that implications and incompatibilities can be offered to play these roles of challenging and defending commitments to accept or reject claimables, and thereby affect the entitlements of those commitments. For giving reasons for and against in this sense, challenging and defending claimables is an integral feature making intelligible the pragmatic normative metaconcepts of commitment and entitlement. The distinctions between statuses of commitment and entitlement, attitudes of acceptance and rejection, reason relations of implication and incompatibility, activities of giving

reasons for and reasons against, and the pragmatic deontic scorekeeping significance of challenging and defending claims must all be understood together. The relations of each to all the others are essential to them being the distinctions they are.

But there is nothing corresponding to giving reasons for and against, challenging and defending assertions and denials, in Restall and Ripley's bilateral normative pragmatics. For that pragmatics retains the fundamentally *monological* character of the proof trees of Gentzen's sequent calculus, which it was designed in the first instance to help us understand. The conception of commitment and entitlement as intelligible only in a context of defending and challenging assertions and denials, of giving and asking for reasons for and against those claims, is by contrast essentially *dialogical*. *I* challenge *your* entitlement to *your* commitment by giving *you* a reason *against* it, and *you* defend *your* entitlement to it by giving *me* a reason *for* it. Reason giving is something that practitioners do beyond simply committing themselves by asserting and denying. It is those further speech acts of giving reasons for and against commitments that can have the pragmatic significance of challenges and defenses, with corresponding normative effects on entitlements to the commitments that are challenged or defended.

The regimented pragmatic metavocabulary I have been using represents reasons theoretically by implications of the form $\Gamma|\sim A$ and $\Gamma\#A$ (staying with the single succedent form for simplicity). We already know how to think about the premises and conclusions of these reason relations. I have only been considering premise sets consisting of sentences expressing doxastic commitments to accept. (The fact that in natural languages we can only assert a set of claims by a list of assertions leaves room, as I've remarked in passing before, for thinking of the premises as forming multisets, or even lists—as Gentzen himself did.) In order to add to claiming a further kind of speech act consisting of giving reasons for or against claims, it is necessary and sufficient to introduce something to do the expressive job of the turnstiles. That is, discursive practices as here construed need something to do the expressive job of the English “therefore,” “so,” or “because.” With such a linking expression, and interlocutor can get the effect of invoking the implication “ $\Gamma|\sim A$ ” as a reason by asserting all of Γ , using the reason-connector “so,” and then *asserting* A , and can get the effect of invoking the incompatibility “ $\Gamma\#A$ ” as a reason by asserting all of Γ , using the reason-connector “so,” and then *denying* A . Helping ourselves to the distinction between asserting the conclusion and denying it recruits the implicit content of the reason relations to make it possible to capture both implication and incompatibility using only the single connector “so” (or “therefore,” or, changing the order, “because”). I called the link-word a “connector” rather than a “connective” because it is not serving to form a new assertible content, as a conditional would do. It does not embed or iterate. Its function is just to connect in a distinctive way speech acts that were already in the picture: assertions and denials.

Enriching the pragmatic metavocabulary by adding the capacity to specify the new compound speech act of giving reasons lets us characterize dialogical discursive practices in which in addition to the basic speech acts of asserting and denying there are speech acts with the pragmatic significance of challenging and defending assertions and denials by giving reasons for and against them. Assertion and denial express commitments to accept or reject. Reasons for and against commitments, used to challenge and defend claims, have the deontic scorekeeping effect of altering *entitlements* to those commitments. In the essay I quoted from above (applauding its understanding of material incompatibility or incoherence as conceptually more basic than negation, in the sense of being prior to it in the order explanation) Ripley contrasts two broad views about the aspects of use that are most important for meaning—within the scope of pragmatics-first approaches to semantics and logic. One, whose avatars include Dummett and Prawitz, looks to relations of *warrant*: of justification and the giving of reasons. The other, which Ripley defends, is focused instead on *coherence*. He is mostly concerned to celebrate the expressive power that the *bilateral* approach he shares with Restall in elaborating the coherence view has over the *unilateral* approach of the warrant theorists. The story I have been telling synthesizes warrant-based approaches, in terms of reasons, and Ripley and Restall’s coherence-based approach within an overarching bilateralist framework. I have just been explaining how doing that requires not only going beyond the unilateralism of previous warrant-based approaches, but to a dialogical construal of discursive practice that goes beyond the monological conceptions that is common to both the warrant-based and coherence-based approaches Ripley surveys.

The principal benefit of making explicit the dialogical structure of discursive practice that incorporates both Restall and Ripley’s bilateralist coherence-based normative pragmatics and practices of giving reasons for and against claims, so rationally challenging and rationally defending them, is to articulate the characteristic defining functional roles played by the key pragmatic metaconcepts of acceptance/rejection, commitment/entitlement, implication/incompatibility, reasons for/against, and rationally challenging/defending claims. A collateral benefit is that it provides what is needed to respond to a fundamental theoretical and conceptual difficulty that becomes visible when one tries to get clear about the meaning of material reason relations of implication and incompatibility that are substructural (in comparison with their purely formal traditional logical relatives) in being nonmonotonic. For nonmonotonicity threatens to undercut one of the principal ways in which I appealed to pragmatic considerations to say what reason relations *are*, what our turnstiles of implication and incompatibility *mean*. It is a manifestation of a conceptual challenge that faces *any* construal of reason relations that does not impose the structural requirement of monotonicity on them.

In Section III I suggested a normative pragmatic understanding of the implication statement “ $\Gamma \vdash \sim A$ ” as making explicit in a pragmatic metavocabulary a deontic scorekeeping commitment effect that commitment to accept all the elements of Γ (its explicit content) precludes entitlement

to reject A (which counts thereby as part of the implicit content of Γ).²⁴ In addition to the distinction between the explicit and the implicit (implied) content of the premise set Γ , which is marked by the implication-codifying snake turnstile, in the pragmatic context of reasoning practices we can ask about the relation between the commitments explicitly included in Γ and other background commitments that an interlocutor might have. Anyone who globally requires that implication relations be structurally monotonic—perhaps because of commitment to logicism about reason relations—has the right to dismiss this question. Collateral commitments are irrelevant to the implication in such a setting, because monotonicity means that *whatever* other commitments one added to Γ , the conclusion A would still follow. In a monotone setting, one can say that commitment to accept all of Γ precludes entitlement to reject A, *tout court*—in the sense of being independent of any further commitments. By contrast, latitudinarian nonmonotonic readings of implication must contend with reasons to reject the pragmatic analysis that take the following form. “In admitting that if we add B to the explicit premises, A no longer follows from Γ together with B, you show that it is just false that commitment to accept all of Γ precludes entitlement to reject A. For accepting all of $\Gamma \cup \{B\}$ includes accepting all of Γ , and yet does not, by your own admission, preclude entitlement to A. The fact that adding commitment to accept B defeats the implication shows that your analysis of “ $\Gamma \vdash A$ ” cannot be right.”

To maintain that $\Gamma, B \vdash A$ is not incompatible (at the metalevel) with $\Gamma \vdash \sim A$, the theorist of nonmonotonic reason relations must read the latter as saying something like “As long as Γ comprises *all* the explicit premises, A follows. That is compatible with admitting that $\Gamma, B \vdash A$, for that is not a claim about what follows from Γ *and that’s all*.” Allowing nonmonotonic implications incurs an obligation to explain the significance of the “and that’s all” clause. Comparison with two other cases might be helpful here. The third proposition of the *Tractatus* is “The world is determined by the facts, and by their being *all* the facts.”²⁵ The status of this last, “and that’s all” condition is a perennial embarrassment. It does not, presumably, state another fact. Nor does anything else in the book evidently help us to understand it. Is it our first indication that the propositions of the book are strictly senseless, “unsinnig”? Again, it has been objected to some substitutional analyses of universal quantifiers that render them as long, perhaps infinite, conjunctions that they must inevitably leave out an important element expressed by universal quantifiers. They are only entailed by conjunctions that somehow also convey that they includes *all* the instances, which it seems no *mere* conjunction can do.

But what does it mean to say that A follows from Γ as long as Γ comprises *all* the premises? It seems we have to say that so that $\Gamma \vdash \sim A$ is not contradicted by observing that $\Gamma, B \vdash A$. The

²⁴ The issue I want to discuss is not sensitive to the difference between single succedent and multisuccedent settings, and has an analogue for incompatibility. So we can stay with the simplest example: single succedent implication.

²⁵ 1.11: “Die Welt ist durch die Tatsachen bestimmt und dadurch, dass es alle Tatsachen sind.”

challenge say what work is done by saying “and that’s all” is quite general. It confronts anyone trying to make sense of nonmonotonic reason relations. But it might seem that it arises in particularly pointed form for the pragmatic analysis endorsed here. For how are we to understand “Commitment to accept all the elements of Γ *and that’s all* precludes entitlement to reject A”? Should it be in terms of counterfactuals of the form “If anyone *were* committed to accept all the elements of Γ *and nothing more*, then they *would be* precluded from being entitled to reject A”? But does that even make sense? Or does it belong in a box with counterfactuals with uninterpretable antecedents such as “If I were a dry well-made match...”? No-one *could* be committed to accept *just* the enumerable elements of any finite set of sentences—and, given the generality of the issue with nonmonotonic material consequences, this includes premise sets restricted to, say, nautical, or culinary, or metereological topics. The context and stage-setting required for anyone to count as having doxastic commitments requires that one have *lots* of them, on many topics (such as other speakers). The subjunctive situation we are being asked to entertain is not just *counterfactual*, it is *counterintelligible*. And the problem with appealing to counterfactuals here is not a consequence of our analysis involving practical attitudes of commitment or acceptance. For it is no help to say instead “If the elements of the premise set Γ were *true*, *and nothing else was*, then A would be *true*. Even if we are allowed to read “nothing else” as “nothing else not implied by them,” (so that the account is not just self-contradictory outside of instances of CO) it is not clear that the antecedent specifies a genuine possibility. (Compare the question of whether it is intelligible to postulate a possible world containing only one object, a human teacher of high-school French.)

If counterfactual analyses of the implicit “and that’s all” stipulation that is needed to make sense of reason relations of the form $\Gamma|\sim A$ in a nonmonotonic context won’t work, what alternatives are there? One suggestion is to invoke a kind of attitude I have not mentioned here: supposition. The thought is that although any interlocutor at any time is committed to accept many (perhaps indefinitely many) claims, and to reject many others, beyond what is or even can be included explicitly in any premise set, they might at any point want to investigate the results of making only a limited number of *suppositions*. The idea is that when we say that Γ nonmonotonically implies A, we are talking about what follows from *supposing* that the elements of Γ are true, and not making any further *suppositions*. Suppositional acceptance of Γ is not a kind of commitment, but something else, perhaps something *sui generis*.

One reason for including some such merely suppositional attitude in an account of reasoning is that, as I claimed earlier, in order to count as grasping claimable contents, interlocutors must have mastered *subjunctive* reason relations. They must have a relatively robust capacity to tell what follows from premises they do not accept, as well as from those that they do accept. My understanding concepts such as copper, lion, and anger involves having at least some views about how this copper penny, the largest lion in the Pittsburgh zoo, and our famously angry neighbor *would* or *could* behave under circumstances I do not believe them

currently to be in—for instance, that the penny would conduct electricity if a current were applied to it, that the lion could eat a small dog, that the neighbor would respond unpleasantly if someone knocked over his garbage can. One cannot count as grasping (having minimally mastered the use of) the content expressed by a sentence if one associates with it only a single implication (or incompatibility). One must have *some* practical know-how in discriminating which of its neighbor implications, with somewhat varied premises and conclusions, also hold. If mastery of reason relations that are subjunctively robust in that they support counterfactual reasoning already involve the capacity to reason from mere suppositions as well as from commitments, then supposition must be part of any account of reasoning and reason relations.

But here, too, we would need an account of what it is to suppose *only* Γ . I might even be committed to accept B, but in asking whether or not $\Gamma \vdash \sim A$, I am only supposing the premises explicitly in Γ , not also B. But supposing only Γ is *not* to pretend that only Γ is true (or to pretend to take-true only Γ). For that does not make sense. So what is it to suppose only Γ ? In the absence of a pragmatic account of what one is doing in supposing Γ rather than accepting it, invoking supposition seems just to relocate the issue without resolving it.

We can clear up these difficulties by looking more closely at the discursive practices within which performances can have the pragmatic significance of undertaking doxastic commitments. The guiding idea of my argument here has been that these are practices of making claims, and defending and challenging them by giving reasons for and against those claims. As I have been using the term, making a claim is undertaking a commitment of a distinctive kind: accepting or rejecting what is expressed by declarative sentences. (Performances count as the utterance of declarative sentences, as specified in a syntactic metavocabulary, just in virtue of playing this role in discursive practice, specified in a pragmatic metavocabulary). Interlocutors' *entitlements* to those doxastic commitments are threatened by the offering of reasons against them, and vindicated by the offering of reasons for them. What count as reasons for and against what is expressed by declarative sentences determine the contents those sentences express, as articulated in a broadly inferentialist semantic metavocabulary of implication and incompatibility.

In order to understand in pragmatic terms what it means for the premises of a reason to be put forward as consisting of the sentences explicitly contained in a finite set Γ *and that's all*, it will suffice to think of extremely simple, stripped-down dialogic practices of undertaking doxastic commitments and assessing interlocutors' entitlements to them—as long as we are careful not to abstract away from any complications that bear on the issue we are interested in.

- We assume that all the participants agree on the reason relations that govern their interactions: that they share a material semantic frame specifying which implications are good and which implications hold.
- We think of commitments as cumulative throughout a dialogue. That is, we think of practices that do not allow retraction of previously undertaken commitments. (“No

backsies.”) Imposing this simplification rules out contrapositive forms of reasoning, which are obviously very important in lots of settings. As pointed out in Section III, omitting these does not at all undercut the distinction between reason relations and reasoning practices, once we understand the former in terms of the constraints they place on joint assignments of commitments and entitlements and the latter in terms of keeping score on these as they fluctuate in the light of new commitments—even though the argument from Harman with which I introduced the distinction does use the example of reasoning contrapositively.

- We take the discursive practice to have a default-and-challenge structure.²⁶ So interlocutors are treated as entitled to their claims (acceptances/rejections) unless and until they are challenged by other interlocutors offering reasons against those commitments. Loss of entitlement due to challenge can be repaired by offering a defense, in the form of a reason for the commitment in question. And both challenges and defenses are treated as successful—removing entitlement and reachieving it, respectively, until and unless *they* are challenged.

So, by contrast to commitments, which are treated as permanent, entitlements fluctuate as reasons are supplied and collide. All entitlements are only provisional and *prima facie*. They are fragile and always in danger of being revoked. For instance, no matter how good one’s reasons are for (say) accepting a claim, one’s entitlement to that commitment can always be lost if one later undertakes a commitment incompatible with it.

- Reasons can be challenged in two ways, either by offering a different reason of countervailing polarity for or against the conclusion of the implication/incompatibility invoked as a reason—offering a reason *against* a conclusion the original reason was a reason *for*, or *vice versa*—or offering a reason *against* one or some of the premises of the original reason.
- Putting forward a reason is a tripartite act, comprising
 - i) Asserting all the sentences in some premise set Γ , so expressing one’s acceptance of them, and
 - ii) Using a reason-expressing connector meaning something like “so” or “therefore” (or “because”, if order of premises and conclusion is reversed). This sort of expression is not a logical connective, because it does not embed or iterate, but connects speech acts.
 - iii) Asserting or denying a conclusion, so expressing one’s acceptance or rejection of it.

The use of the connector between endorsing the premises and taking up a stance toward the conclusion marks the fact that a reason is being offered, in the form of the

²⁶ I talk about the nature and importance this sort of practical normative structure has for epistemology in Chapter 3 of *Making It Explicit*. Shahid Rahman’s group (see next note) also endorses this fundamental structure of dialogues. They say “The Oslo and the Stockholm lectures of Martin-Löf (2017) contain challenging and deep insights in dialogical logic, and **the understanding of defences as duties and challenges as rights** is indeed at the core of the deontics underlying the dialogical framework.” (p. 293)

implication of A by Γ or the incompatibility of A with Γ . (Recall that our first simplifying assumption is that there is universal agreement concerning the reason relations that articulate the contents expressed by declarative sentences.) If $\Gamma|\sim A$, then accepting all of Γ provides a reason *against* rejecting A and a reason *for* accepting A, because it precludes entitlement to reject A. If $\Gamma\#A$, then accepting all of Γ provides a reason *for* rejecting A and a reason *against* accepting A, because it precludes entitlement to accept A.

That reasons must be finitely statable to be usable to challenge or vindicate entitlement to claims in dialogue at all is not a simplifying assumption. It is an essential feature of discursive practices. To put forward a set of premises as a reason for one's own commitment or against another's is to put it forward as having the pragmatic significance of altering the entitlements that go with various commitments: as reclaiming an entitlement for oneself or as removing an entitlement from another. Part of what lies behind the "and that's all" is that only what is in Γ , and nothing more, is needed to provide a reason with the desired effect on entitlements: for it to serve as a defense or as a challenge. For every claim I make, including those made as endorsements of the premises of reasons I offer, whether in mounting a rational defense or a rational challenge, makes me vulnerable to challenge in turn. Since commitments come with that practical cost, they are not to be undertaken beyond necessity. Of course, in nonmonotonic contexts, one way to object will be to cite a contrary reason whose premises are a superset of the one originally offered. But that possibility doesn't make the original move useless, or undercut its significance at all—any more than the possibility at that point in the dialogue of offering a countervailing reason of any other shape does.

Another perspective on the features of dialogic discursive practice that make the "and that's all" intelligible is provided by what Shahid Rahman calls the "Socratic Principle."²⁷ This is the idea that:

- Any commitment another interlocutor explicitly undertakes is available to anyone to use, as a premise both in arguing against that interlocutor's entitlements and in defending one's own, insulated from challenge by that interlocutor.²⁸

Once put on the conversational table, those commitments become a public resource. While the Socratic Principle is important in making possible the securing of common ground—commitments any interlocutor can appeal to—putting premises forward as potential common

²⁷ Shahid Rahman, Zoe McConaughy, Ansten Klev, Nicolas Clerbout *Immanent Reasoning or Equality in Action: A Plaidoyer for the Play Level* [Springer 2018].

²⁸ Perhaps the Socratic Principle is where we should look to find the essentially dialogic pragmatic basis for understanding what one is *doing* in merely *supposing* that Γ . But I will not pursue this thought here. Rahman does, saying in the Preface to the book just cited (p. vii): "...the present book can indeed be read as furthering Sundholm's own extension to inference of Austin's remark on assertion acts; Sundholm did indeed produce this forceful formulation: "When I say 'therefore', I give others *my* authority for asserting the conclusion, given *theirs* for asserting the premisses."

ground also provides further resources for an interlocutor's critics.²⁹ Part of what one is doing in putting forward Γ "and that's all" as a reason entitling one to a commitment or challenging another's commitment is offering to the other just those premises as "free moves"—not themselves subject to challenge or in need of defense—and no more (in this speech act).

We should understand what is invoked by the "and that's all" demarcation of premise-sets in terms of all these features of the pragmatic significance of the distinction between what is and what is not in the premise set Γ that an interlocutor puts forward as a part of a reason in the form of an implication or an incompatibility, distinguishing them from the interlocutor's other commitments, avowed or not. And the point I have been leading up to is that those features that explain the significance of the limitation of explicit premises of any and every reason arise from very basic features of the essentially social, dialogical character of discursive practice: the way reasons challenging and reasons defending commitments are traded, undertaking further commitments and collectively determining the ebb and flow of entitlements to all those commitments. As such, the salient normative significance of the limited premise sets of all the reasons exchanged even in dialogues with the simplified structure I sketched are essentially independent of whether the implications and incompatibilities playing the functional, entitlement-altering role of reasons are monotonic or not. What I have been calling the "range of subjunctive robustness" of implications and incompatibilities (codified as the Υ -sets by which they are semantically interpreted in Kaplan's implicational phase-space semantics) is the set of additional premises that could be added to the original premise set Γ while preserving the goodness of the reason relation, and so the pragmatic potential of the reason to confer or remove entitlements. Imposing a global structural requirement of monotonicity is insisting that ranges of subjunctive robustness must be *total* in order for the reasons they are associated with to be intelligible *as* reasons. But in fact we can see that the structural difference between monotonic and nonmonotonic reason relations makes very little practical difference to the pragmatic significance and dialogical tractability of reasons.

The key point, I think, is that even implications or incompatibilities that are not *semantically* defeasible by the addition of further collateral premises or auxiliary hypotheses are still *pragmatically* defeasible in the sense that the entitlements they confer can be defeated by contrary reasons. The *only* practical difference is that an additional kind of objection to a reason such as $\Gamma|\sim A$ comes into play in a nonmonotonic setting. In a monotonic setting it is possible to accept the goodness of an implication $\Gamma|\sim A$ but not take that to preclude entitlement to reject A because one asserts some further set of premises that are incompatible with one of the elements of Γ . Removing that entitlement voids the defense of A that the implication would otherwise afford. In a nonmonotonic setting there is another way to achieve the same effect. One can

²⁹ I report some research we did (implemented in computer programs) on simple dialogues with this shape in "The Pragmatist Roots and Some Expressivist Extensions of *The Dialogical Roots of Deduction*" written for a symposium on Caterina Dutilh Novaes' book of that name. (It is available at my website: www.pitt.edu/~rbrandom.)

accept all of Γ , conceding entitlement to those commitments, and accept the implication, but mount an entitlement-voiding challenge to A by asserting some further claim B such that $\Gamma, B \sim A$ or even $\Gamma, B \# A$. Such a critic can accept all of Γ and not be precluded from entitlement to reject A even though (as all parties agree) $\Gamma \sim A$. (Γ thereby becomes part of the common ground, though A is controverted.) And the possibility of such a riposte by a critic does not make it pointless for the claimant to have offered the reason $\Gamma \sim A$, any more than in the case of *de jure* indefeasible implications the possibility of accepting the implication but offering reasons to reject one of its premises makes it pointless to offer indefeasible implications or incompatibilities as reasons. For the entitlements offered by reasons are *always* vulnerable, for instance to being precluded by collateral incompatible commitments, and in *that* sense defeasible, even when the reason relations themselves are *indefeasible*. The fact that pragmatic defeat is always in principle available means that it makes no structural *pragmatic* difference to the way dialogues work whether or not *semantic* defeat (nonmonotonicity of implication/incompatibility) is available.

In a pragmatics-first order of explanation, we are to understand the significance of a limited, finitely specifiable set of premises in terms of the pragmatic significance putting forward reasons with such premise sets has in the context of practices of reasoning: making claims and challenging and defending them with reasons. Offering a reason against another's commitment (acceptance or rejection) removes its entitlement, and offering a reason for one's own commitment restores its entitlement. In each case, the premises of the reason offered are further commitments whose entitlement can be queried. One opens oneself up to that sort of deontic vulnerability only with the premises that are explicitly included in the premise set. The only difference the nonmonotonicity of (most) material implications and incompatibilities makes is that entitlement to one's conclusions can also be challenged by other interlocutors who offer reasons whose premise sets are supersets of one's own, and are therefore, by the Socratic Principle, not subject to further challenge by anyone who already endorsed them. In this way we can understand the feature of nonmonotonic reason relations that is acknowledged by an "and that's all" clause in terms of the practice in which reason relations are invoked to alter entitlements to commitments that have been explicitly avowed.

This is a special case of a more general pragmatics-first explanatory strategy. The basic point is that reason relations of implication and incompatibility themselves—what we express in our metalanguage by the two turnstiles, ' \sim ' and ' $\#$ '—are to be understood in terms of the role reasons citing them play in discursive (reasoning) practices of challenging and defending doxastic commitments: the acceptance or rejection of claimables. Indeed, even those claimables themselves, the 'propositions' expressed by declarative sentences, what can be accepted or rejected, are to be understood in terms of their roles in discursive, reason-giving practices, as crystallized in the reason relations of implication and incompatibility that normatively govern those practices, determining the significance of acceptances and rejections. This thought is

operationalized in Kaplan's implicational phase-space semantics, where propositions are represented by pairs of \forall -sets: the premissory and conclusory roles of a sentence as what must be adjoined to the sentence, on either side of the turnstile, to make good implications.

VIII. Conclusion

Two overarching aims have guided the story told here. The first is to sketch an order of explanation that begins in the pragmatics of discursive practice: in an account of what it is one is *doing* in claiming and reasoning. Semantics then takes the form of an account of the content conferred on acts and practical attitudes by the roles they play in the discursive practice that has been specified in a pragmatic metavocabulary. The distinctive pragmatic expressive role played by logic, and the sort of semantic content logical expressions and concepts acquire by playing that role, are then explained in terms of the more general approach to pragmatics and semantics that has been put in place in this way of working out a “pragmatics-first” order of explanation. This way of understanding things contrasts with what I take to be the more traditional and familiar “semantics-first” order of explanation. That contrary methodological approach starts with the contrast between truth and falsity, and understands the truth-evaluable contents expressed by declarative sentences in terms of these “truth values”. The genus of semantics-first explanations includes as species truth-conditional theories (such as Davidson’s), and accounts (such as Lewis’s) that begin with extensions and proceed to define semantic intensions as functions from indices (such as possible worlds) to extensions. The version I have addressed here is Kit Fine’s truth-maker semantics, which stands out as the premiere example of the semantics-first approach because of its technical sophistication and flexibility, and its conceptual adaptability to and fecundity in addressing a wide variety of topics and issues. In addition, it is congenial to the line of thought pursued here in offering a general account of material, nonlogical content that is not inspired by or adapted from a semantics for specifically logical vocabulary (not even modal vocabulary). The semantics it offers for logical vocabulary shows up as just one more special case. I suppose that many who adopt a semantics-first order of explanation (whether utilizing the truth-maker semantics or not) will envisage building a Stalnaker-type pragmatics on top of it.

The second orienting aim is to articulate the lessons we can learn about the *structure* of reasoning and the reason relations that normatively govern it, by adopting the pragmatics-first direction of explanation. Here the relevant contrast is with a “logic-first” order of explanation of the structure of reasons. That way of proceeding models its understanding of implication in general on specifically *logical* implication. Indeed, in the most extreme case, logicism about reason relations, it *identifies* implication relations generally with implication relations that hold in virtue of logical vocabulary alone (in the sense of being good under arbitrary substitutions of nonlogical for nonlogical vocabulary). One consequence of adopting this logistical approach is not appreciating that reason relations of material *incompatibility* are coeval and co-ordinate with reason relations of *implication*. For focusing on logical reason relations encourages either mistaking the significance of notational convenience, when incompatibilities are assimilated to

implication relations by the device of empty right-hand sides in Gentzen's proof-theoretic notation, or leaving incompatibility to be handled downstream of the introduction of logical negation, as in Tarskian model-theoretic presentations. As a result, the essentially *symmetric* structure of incompatibility in general has not been much remarked on or thought about. Further, the substantial structural differences between (in general nonsymmetric) nonlogical, material relations of implication and logical implication have not been sufficiently appreciated. The nonmonotonicity (the failure of the structural principle MO, that a good implication remains good upon arbitrary addition of further explicit premises) of nonlogical reasoning in general is widely acknowledged. But the dominant strategy of nonmonotonic *logics* has been to try and reconstruct it by building on top of *de jure* monotonic logical implications. The failure of nonlogical implications in general to remain good upon explicitation of their own consequences—the failure of even “cautious” monotonicity (CM)—has not been the subject of much theoretical investigation. And the fact that in the nonlogical case implications are not in general even transitive (that CT fails for them) has also not been a well-known phenomenon. (Failures of CT have principally been of interest to those concerned to address semantic paradoxes.) These substantial structural differences between reason relations in general and specifically *logical* reason relations put major constraints on both semantic and logical codifications of these relations of implication and incompatibility for the general case.

One bad effect of logicist metaconceptual oppression has been rendering largely invisible the important intellectual operation of *explicitation*, in the sense of making explicit as further premises conclusions that before the operation remained merely implicit as implied conclusions. CM and CT together imply the *inconsequentiality* of explicitation: that it can in principle never make any difference to what a premise set implies. In many ordinary cases this is simply not so. Explicitation in this sense deserves more epistemological investigation than it has traditionally had.

I began exploring the first theme by observing that there is a fundamental bipolarity that shows up in both semantic and pragmatic terms: as the distinction between the truth values true and false and between the practical attitudes of accepting and rejecting, *taking-true* and *taking-false*, expressible by speech acts of assertion and denial. One might doubt that, in the end, one can understand either of these without understanding the other. (That is, on might take them to be reciprocally sense-dependent conceptions). But I elaborate a pragmatics-first order of explanation that asserts a conceptual asymmetry. A semantics-first order of explanation proceeds to understand the bearers of truth values, what can in the first instance be true or false (what is, accordingly, expressed by declarative sentences) in terms of truth or falsity *conditions*—or, in the best version currently available, truth-makers and falsifiers. In the pragmatics-first order of explanation, I argued, we begin with the idea that it is an integral feature of making claims, undertaking doxastic commitments, whether by accepting or rejecting something expressed by a declarative sentence, that one is liable to *challenge* by others making

claims that provide reasons *against* one's commitment, and obliged to *defend* one's commitment by making further claims that provide reasons *for* that commitment. Implicit in these rational *practices*, practices of giving reasons *to* do something doxastically (namely, accept or reject), we find two sorts of reason *relations*: *implication*, articulating reasons *for* what can be accepted or rejected, and *incompatibility*, articulating reasons *against* what can be accepted or rejected.

At this point, as expressed in the pragmatic metavocabulary put in place for talking about what interlocutors *do* when engaging in discursive practices, we have two kinds of *commitment*: acceptance and rejection. In addition to commitments, we have the kind of *entitlement* that is altered by the giving of *reasons*. Commitments and entitlements differ in their pragmatic structure, in that commitments relatively durable. Even if we do not think of them as irrevocable, they can always be reinstated by adopting the appropriate attitude of acceptance or rejection, paradigmatically by asserting or denying something. By contrast, entitlements are fragile and temporary, and not under the control of the subject in the same way that commitments, at base, are. Accompanying commitments by default, in the most basic case, they admit both removal by challenging reasons and reinstatement by defending reasons. From the pragmatic significance other claims can have as rationally challenging or defending a commitment we understand being a reason *to* accept or reject, and so the reason *relations* of being a reason for (implying) or being a reason *against* (being incompatible with).

The aspiration at the core of this order of explanation is then to move from pragmatics (studying use or practice) to semantics (studying meaning or content) by understanding the claimables expressed by the declarative sentences that are asserted or denied, what can be doxastically accepted or rejected, in terms of the functional roles sentences play in reason relations of implication and incompatibility. Consideration of reasoning practices leads to the overarching distinction between such practices and the reason relations of implication and incompatibility—and so to the distinction between the topics of pragmatics and semantics. The paired basic structural distinctions on the side of reasoning practices are those between the doxastic *attitudes* of acceptance and rejection and the normative *statuses* of commitment and entitlement. Following Simonelli's synthesis of the bilateralisms of Restall and Ripley, on the one hand, and Smiley and Rumfitt, on the other, we can deploy these paired distinctions to offer parallel pragmatic accounts of the reason relations. To say that Γ *implies* Δ is in dialogical deontic scorekeeping terms to say that *commitment to accept* all of Γ precludes *entitlement to reject* all of Δ . To say that Γ is *incompatible* with Δ is in dialogical deontic scorekeeping terms to say that *commitment to accept* all of Γ precludes *entitlement to accept* all of Δ . This reconstruction is in the spirit of the other bilateralists, but where Restall and Ripley appeal to a single, holistic, undifferentiated notion of incoherence (being normatively "out of bounds"), we show how the distinction between the two deontic statuses of commitment and entitlement can be deployed in concert with the bilateralist dyad of acceptance and rejection to yield a pragmatic reading of the two different (but intimately related) semantogenic reason relations of implication

and incompatibility. This is adding substantial pragmatic fine structure beyond what previous bilateralists have offered.

We saw in these terms how to understand a *pragmatic* concept of implicit commitment. If Γ implies Δ , then commitment to accept all of Γ *explicitly* precludes entitlement to reject all of Δ and it *implicitly commits* one to accept something in Δ . If Γ is incompatible with Δ , then commitment to accept all of Γ *explicitly* precludes entitlement to accept all of Δ and it *implicitly commits* one to *reject* something in Δ . This pragmatic concept of further commitments that are *implicit* in commitment to accept a premise set (by contrast to the entitlements explicitly precluded thereby) is then echoed later on by a semantic concept of the implicit content of such premise sets. These are the further sentences that are implicit in the sense of implied (or exclude d) by it. Moving a sentence from the right-hand side, marking the consequences of an implication, to the left-hand side, marking its premises, is accordingly a kind of *explicitation*: turning implicit (implied) content into explicit premises of an implication. This semantic sense of “making explicit what is implicit” played a central role in one of the principal arguments I offered against extending the structural principles of Cautious Monotonicity (and so, the stronger principle of Monotonicity) and Cumulative Transitivity (Gentzen’s “Cut”) from the realm of *logical* consequence relations to *prelogical*, *material* consequence relations.

The metasemantic project of understanding the contents expressed by declarative sentences—what one can be committed or entitled to accept or reject—in terms of the roles they play in prelogical relations of implication and incompatibility goes hand in hand with the aim of extracting the most fundamental *structural* features of the reason relations that articulate reasoning practices from an account couched in this pragmatic metavocabulary. These are the features that should be held constant if one Ramsifies the pragmatic story, putting variables in place of its key concepts: acceptance/rejection, commitment/entitlement, implication/incompatibility, premise set/conclusion (set), declarative sentence/claimable content, challenge/defense. Such a Ramsified theory then codifies the functional roles that items must play in social practices for those practices to qualify as *reasoning* practices, for those relations to qualify as *reason* relations, and for the various normative statuses, acts, and attitudes to count as *discursive*, *doxastic* statuses, acts, and attitudes. What such a functionalist theory codifies is the practices of deontic scorekeeping on commitments and entitlements that evolve throughout a dialogue that proceeds by claiming, and rationally challenging and defending those claims. This is what it is for a community of interlocutors in practice to take or treat some marks or noises (Sellarsian “sign designs”) as sentences expressing attitudes of acceptance and rejection, which count as undertaking *doxastic commitments* in virtue of the role they play in practices that count as *discursive* in virtue of the role *reason relations* of implication and incompatibility play in them in *rationally challenging* and *defending* those commitments, thereby systematically affecting their corresponding *entitlements*.

We saw that there is a striking structural asymmetry between the two reason relations. Incompatibility is *de jure* symmetric, in the sense that if, possibly in the context of some set Γ of collateral premises serving as auxiliary hypotheses, A is incompatible with B, then B is in the same context incompatible with A. By contrast, if in the context of Γ , A implies B, it does not follow that in that context B implies A. In the absence of methodological commitment to a logic-first order of explanation, which understands material incompatibility (Aristotelian contrariety) in terms of logical inconsistency (Aristotelian contradiction), it is not initially clear why this should be. Thinking in terms of implicit pragmatic commitments, if in the context of commitment to accept all of Γ , commitment to accept A implicitly commits one to reject B, why, in the same context, should commitment to accept B implicitly commit one to reject A? Thinking in terms of the explicit significance of incompatibility relations helps some, for $\Gamma, A \# B$ can be read as telling us that one cannot be entitled to accept all of Γ , A, and B. And that is a symmetric relation. But it is really only when (again following Simonelli) we look at the dialogic pragmatics of using incompatibilities to keep score on the entitlements of others that a Dutch-Book-like argument becomes visible for the necessary symmetry of incompatibility.

Kaplan's substructural implicational phase-space semantics shows in detail how an expressively powerful formal semantics can be elaborated from the material relations of implication and incompatibility that precipitate out of the functionalist story told in such a normative pragmatic metavocabulary. Declarative sentences are understood as what can appear as either premise or conclusion, what can play roles on both sides of relations of implication and incompatibility. The contents expressed by declarative sentences are then the roles they play in the two sorts of reason relation. I showed that this implicational phase-space semantic metavocabulary has the same expressive power as Fine's truth-maker semantic metavocabulary, in the sense that the reason relations of implication and incompatibility expressible in the two frameworks can be matched up to isomorphism. In this respect and to this extent, that is, insofar as relations of implication and incompatibility are concerned, the semantics-first truth-maker semantic metavocabulary and the pragmatics-first implicational semantic metavocabulary are on a par.³⁰

That result is a corollary of Hlobil's remarkable demonstration that, with a very natural definition of consequence in the truth-maker framework— $\Gamma | \sim \Delta$ iff the fusion of every verifier of all of Γ with a falsifier of all of Δ is an impossible state—reason relations in the truth-maker framework can be understood according to the bilateral normative pragmatic reading Restall and Ripley recommend. Their bilateralism takes " $\Gamma | \sim \Delta$ " to say that the position that consists of accepting all of Γ and rejecting all of Δ is normatively "out of bounds." I have suggested that that status of "out of boundness" is in turn naturally rendered in terms of preclusion of joint

³⁰ The result depends on using the consequence relation Hlobil crafted in the truth-maker framework to mirror Restall and Ripley's bilateral normative pragmatic reading of the multisuccedent turnstile. Fine himself does not use that version of implication, preferring a variety of others, which Hlobil locates with respect to his preferred one.

entitlement to all of those commitments (acceptances and rejections). Hlobil's result accordingly shows that the normative deontic scorekeeping language that the pragmatics-first order of explanation uses to explain the reason relations of implication and incompatibility is an expressively adequate pragmatic metavocabulary for Fine's truth-maker semantics. It is also an expressively adequate pragmatic vocabulary for Kaplan's implicational phase-space semantics. (The previously cited result shows that what works to codify the reason relations of the one semantic framework will work to codify the reason relations of the other.) The result is that a pragmatics-first order of explanation that employs a suitable normative pragmatic metavocabulary not only suffices to say what one is doing in using a semantic metavocabulary, but applies to and can explain the use of both broad species of semantic metavocabulary: the truth-maker semantics, which is the most sophisticated development of the semantics-first tradition and the implicational phase-space semantics, which is the most natural and proximal pragmatics-first semantics.

So the very same normative pragmatic metavocabulary—of commitments to accept or to reject, and of preclusion of entitlement to such commitments—that can be used to specify the reasoning practices in which sentences of the prelogical object language are used to make, challenge, and defend claims, *also* suffices to specify the use of the semantic metavocabularies to characterize both reason relations of implication and incompatibility and what is expressed by the declarative sentences that can be accepted or rejected, true or false. This result of high theory—concerning pragmatic metavocabularies *for (inter alia)* semantic metavocabularies—provides good reasons for an explanatory privileging of pragmatic over semantic metavocabularies, specifically, for a pragmatics-first order of explanation. It also offers a new argumentative route entitling us to identify what is truth-evaluable as true or false with that to which one can take the practical normative attitudes of accepting or rejecting. It is the very same contents that can *be* true or false that can practically be *taken*-true (accepted) or *taken*-false (rejected). Contents suitable for both truth-evaluation and doxastic commitment can be specified in terms of the role they play in reason relations of implication and incompatibility. And that is so whether or not we understand those reason relations immediately in pragmatic terms, and then construct a directly implicational semantics for them, or take them as reflecting antecedently intelligible conceptual contents taking the form of pairs of sets of verifiers and falsifiers.

Indeed, the fact that one can use the same normative pragmatic metavocabulary that makes it possible to say what one is doing in using ordinary, ground-level nonsemantic, nonlogical vocabulary to understand what one is doing in using the truth-maker semantic metavocabulary to codify the reason relations that propositional contents stand in to one another reveals an underlying *asymmetry* in the distinction between truth-evaluations of claimables as true or false and practical attitudes of accepting or rejecting, on which the distinction between semantics-first and pragmatics-first orders of explanation is founded. For the pragmatic metavocabulary for the truth-first semantic metavocabulary underlines the fact that what the semantic theorist is *doing* in

sorting or evaluating claimables to begin with as true or false (perhaps guided by a view about what states would verify or falsify them) is just what the pragmatic metavocabulary takes as adopting the basic practical doxastic attitudes: *taking-true* (accepting) and *taking-false* (rejecting). The pragmatics-first order of explanation begins by explicitly theorizing about those practical attitudes as they show up in the use of the object language. The semantics-first order of explanation begins by practically adopting such attitudes, implicitly, and in an untheorized way, as part of the unexplained, taken-for-granted use of its semantic metavocabulary. The attitudes are fundamental in either case. The difference is just how theoretically and methodologically self-conscious one is about them. In the semantics-first order of explanation, the issue of what one is doing in making truth evaluations in the semantic metavocabulary, and in particular, what reasons entitle one to privilege *these* takings-true and takings-false (acceptances and rejections) is resolutely kept off-stage. This seems a point in favor of the pragmatics-first approach.

There is a complementary methodological complaint that might be forwarded to challenge this assessment. One might wonder what it is that qualifies the “reason relations” treated as such by some community of interlocutors as genuine relations of *rationally* following from or ruling out (including or excluding). What if what the community *takes* as following from a particular premise set does not *really* follow from it? I have rejected the logicist conception of reason relations in general as holding in virtue of underlying *logically* good relations of implication and incompatibility—the idea that “good reason” means “logically good reason”—diagnosing it as resulting from a fundamental misunderstanding of the expressive role characteristic of logical vocabulary. But one might have other substantive views about the nature of genuine rationality, perhaps even ones also rooted in pragmatic considerations, such as the view of rational choice theory, which appeals to utility maximization. Or one might prefer a Bayesian framework that might not appeal to preference and utility. In any case, is it not possible that a community whose practices have the structure outlined here should still be radically wrong about what is a genuine reason for and against what? In the end, isn’t the pragmatic functionalist account of reasons offered here just a version of the disreputable Rortyan pragmatist view that “A good reason is whatever your community will let you get away with (will treat as such),” supplemented by a fuller story about what “treating something as a good reason” consists in?

It is important to realize that such an objection presupposes a notion of the semantic contents of claimables that is intelligible antecedently to and independently of the practices of reasoning with them out of which reason relations of implication and incompatibility are precipitated according to the pragmatics-first order of explanation. Rational choice theory assumes that the meaning of the options and the outcomes is fixed in advance of thinking about rationality. Bayesian programs assume that we already understand the claimable contents to which we attach prior and posterior credences and conditional probabilities, in advance of making these assessments of them. When one worries about whether what the community takes to be good reasons in general *are* good reasons, one is ignoring the Quinean lesson that the meanings and

beliefs expressed by their manipulations of sign designs are not fixed independently of the use they make of them. It is of the essence of the view articulated here to see reason relations and meanings as fixed together (by the pragmatics), neither one being intelligible except in a context that includes the other. The structure of reason relations is the structure of propositional meaning (as a proper semantics will tell us), and it is determined by the structure of dialogical discursive reasoning practices (as a proper pragmatics will tell us). It is being used the way they are, playing the roles they do in implications and incompatibilities, that determines what contents are expressed by declarative sentences.

Returning to the issue of the structure of reasoning relations as it appears from a pragmatics-first methodological perspective, logic comes into the story I have told twice: in a traditional role as providing a bad structural model of prelogical implication and incompatibility relations, and in a progressive expressivist version. As to the first, there has been a regrettable Procrustean tendency to impose on implication and incompatibility generally, extraneous structural conditions that are imported from specifically *logical* implication and incompatibility, as originally construed by Tarski and Gentzen. Logical reason relations were the first ones we learned to codify, both in logic, and in semantics, and then further in the truth-based semantics offered for logical vocabulary. As a result, even philosophers who do not fall into the error of logicism about reason relations generally have been tempted to think of prelogical implication and incompatibility relations as structurally monotonic, or, failing that, at least as obeying cautious monotonicity (being subjunctively robust under the addition of implied consequences), and as *de jure* transitive. In fact none of these structural conditions holds of material reason relations in general. Traditional logics are accordingly impotent to express those relations of implication and incompatibility.

The expressivist view of the task of logic I have been recommending understands logic as having a semantic intent. The expressive role distinctive of logical vocabulary is to make explicit, in a material, nonlogical base vocabulary extended by the addition of logical vocabulary, the reason relations of implication and incompatibility in virtue of which the sentences of that base vocabulary mean what they mean (have the contents they do). This is a kind of semantic descent, since the semantogenic reason relations are to be expressed in an extended object language, rather than in a separate semantic metalanguage. The challenge is to make our logic and semantic metavocabularies sufficiently expressively powerful and flexible to do justice to the structure of actual reason relations—what is, from the point of view of traditional logic, their radically *substructural* character. The Ketonen connective rules of the nonmonotonic multisuccedent (NM-MS) variant of the logic G3cp that I rehearse in Section VI (which in a fully Tarski-Gentzen structural setting is equivalent to classical logic, in the sense of determining the same logical theorems and consequence relations) have been shown by Kaplan to find a sound and complete expression in the implicational phase-space semantic metavocabulary, even in the radically substructural case where MO, CM, and CT fail to hold in

general of the base language extended by the addition of conditionals codifying implications and negation codifying incompatibilities. The purely logical reason relations determined thereby are fully structural. But the full implication and incompatibility relations of the conservatively extended nonmonotonic, intransitive object language remain substructural. It is a logic adequate for expressing nonmonotonic, nontransitive reason relations rather than a traditional nonmonotonic logic.

Both the semantics-first truth-maker semantic metavocabulary and the pragmatics-first implicational phase-space semantic metavocabulary are expressively powerful enough to specify the contents expressed by logical vocabulary, including NM-MS. That remains true even if the underlying material consequence and incompatibility relations on the basis of which the logical connectives are introduced are radically substructural. The implicational phase-space semantics was built with exactly this expressive capacity in mind. And the isomorphism of reason relations specifiable in the two semantic frameworks, demonstrated above, shows that the truth-maker semantic framework can be given that extended expressive capacity, too, by using the definition of implication in terms of truth-makers that Hlobil suggested in order to bring out the parallel with Restall and Ripley's bilateralism.

So a pragmatics-first order of explanation enables revealing new approaches both to semantics and to logic. Along the way, it illuminates important features of the structure of reasoning practices and reason relations, and of the relations between dialogic pragmatics and broadly inferential semantics. Furthermore, both the semantic metavocabulary and the logical vocabulary discussed here are expressively powerful enough to capture, each in its own way, the radically substructural relations of implication and incompatibility that are actually exhibited by our use of ordinary, nonlogical, nonsemantic vocabulary, and that articulate the meanings or contents expressed by that vocabulary. The normative dialogic pragmatic metavocabulary deployed here suffices to specify both the use and the meaning-articulating reason relations governing not only substructural material base languages, but also both kinds of semantic metavocabularies we considered—truthmaker semantics and implicational phase-space semantics—and the logical vocabularies presented here as ways of codifying the reason relations of those substructural presemantic, prelogical base vocabularies.

End

[37,105 words]