August 14, 2022

Outline of *Reasons for Logic, Logic for Reasons*

Introduction:
[Brandom]

I. Logic and Reasoning

The beating heart of this book is a distinctive approach to the philosophy of logic: logical expressivism. It articulates an answer to the central, orienting philosophical question about logic: how to understand the relations between logic and right reasoning. The normative center of reasoning is the practice of assessing reasons for and against conclusions. Reasons *for* conclusions are normatively governed by relations of *consequence* or *implication*. Reasons *against* conclusions are normatively governed by relations of *incompatibility*. These relations of implication and incompatibility, which constrain normative assessment of giving reasons for and against claims—the essential critical practices of defending and challenging commitments—are the first significant level of structure of reasoning. These positive and negative reason relations are the aspect of reasoning practices that logic directly addresses.

Understanding the topic of logic in terms of reason relations of consequence and incompatibility (conceptual inclusion and exclusion) goes back as far as the origin of the subject in Aristotle. He classified the forms of judgements, in effect, by the dual functional roles they play in the different figures of syllogism and in the square of opposition—that is, their roles in his codification of relations of consequence and incompatibility. A different note was struck at the dawn of modern logic, however. Michael Dummett says:

> ...[I]n this respect (and in this respect alone) Frege's new approach to logic was retrograde. He characterized logic by saying that, while all sciences have truth as their goal, in logic truth is not merely the goal, but the object of study. The traditional answer to the question what the subject-matter of logic is, however, that it is, not truth, but inference, or, more properly, the relation of logical consequence. This was the received opinion all through the doldrums of logic, until the subject was revitalized by Frege; and it is, surely, the correct view.[1]

And

[1] Michael Dummett, *Frege's Philosophy of Language*: Harper & Row, New York, 1973 (hereafter *FPL*), p. 432.

> It remains that the representation of logic as concerned with a characteristic of sentences, truth, rather than of transitions from sentences to sentences, had highly deleterious effects both in logic and in philosophy. In philosophy it led to a concentration on logical truth and its generalization, analytic truth, as the problematic notions, rather than on the notion of a statement's being a deductive consequence of other statements, and hence to solutions involving a distinction between two supposedly utterly different kinds of truth, analytic truth and contingent truth, which would have appeared preposterous and irrelevant if the central problem had from the start been taken to be that of the character of the relation of deductive consequence.[2]

In these passages Dummett exhibits the common bad habit of logicians of suppressing discussion of incompatibility and focusing exclusively on consequence. This bias is connived at by notational conveniences (particularly in sequent calculi) but it, too, has had deleterious philosophical consequences—in this case obscuring the expressive role characteristic of negation.

The need to distinguish clearly between reason relations and reasoning practices, in spite of their mutual presupposition, is underscored by Gilbert Harman.[3] In keeping with the habitual practice just noted, he argues that we must not confuse *implication* with *inference*, by arguing for the intentionally provocative claim that "there is no such thing as rules of deductive inference." If there were such logical rules, presumably a paradigmatic one would be: If you believe *p* and you believe *if p then q*, then you should believe *q*. But that would be a terrible rule. You might have much better reasons against *q* than you have for either of the premises. In that case, you should give up one of them. He concludes that we should distinguish *relations* of *implication*, from *activities* of *inferring*. The fact that *p, if p then q*, and *not-q* are incompatible, because *p* and *if p then q* stand in the implication relation to *q*, normatively *constrains* our reasoning activity, but does not by itself *determine* what it is correct or incorrect to do. (Notice how reasons against and incompatibilities creep in, even though they are not explicitly thematized alongside implication.)

The result of these considerations is that the relations between logic and reasoning practices are mediated by reason relations, of which at least the two principal species are implication and incompatibility, structuring reasons for and reasons against, respectively.[4] For most of the twentieth century an implicit consensus prevailed as to the answer to the question of how logic relates to reasoning: logic provides a canon of good reasoning, in the sense of determining the reason relations that govern it. This view might be called "logicism about reason relations." In its simplest and purest form, it holds that good reasons just are *logically*

---

[2] Dummett, *FPL* p. 433.

[3] [ref.]

[4] Having acknowledged the importance of distinguishing reason relations from reasoning practices, and so, implication (and incompatibility) from inference, I will risk talking about both implication and incompatibility as "inferential" relations.

good reasons. Logic is what makes good inferences good, in the sense that behind every correct inference there is a logically valid implication (or a logical inconsistency, for inferences leading to the rejection of a claim). If an argument is not underwritten by a valid logical form, then its conclusion does not really follow from its premises.

It came to be acknowledged that it might not be possible to hold all theoretical reasoning to this strict logical paradigm. The central deductive business district might be surrounded by suburbs of inductive or abductive reasoning. Perhaps some analogue of the norms of deductive logic might be found for these less well-regulated outlying regions. Or perhaps they involve different senses of "good reason," which are in various ways parasitic on the paradigmatic logical sense. The general animating logicist thought remained that logic provides not only the paradigm of good reasoning, but also norms constraining reasoning in general. Even if it is not always a sufficient guide (cannot be relied upon in all cases to dictate what conclusions should be drawn), logic sets necessary boundaries beyond which reasoning ought not, and correct reasoning cannot, stray.[5]

Logical expressivism claims that logicism gets things backwards. Logic does not provide a substantive standard for right reasoning in the sense of dictating the correct reason relations of implication and incompatibility: what really follows from or rules out what. It provides expressive tools that let practitioners make explicit the inferential (implicational and incompatibility) commitments that are implicit in their reasoning practices—whatever those commitments are. Here "making explicit" means "putting in the thinkable, assertible form expressed by declarative sentences." What is explicitly expressed by declarative sentences can in turn be understood as what can both serve as and stand in need of reasons: what can play the role both of premise and of conclusion in inferential relations. Prelogical vocabulary lets us make doxastic commitments explicit. Logical vocabulary lets us make explicit inferential commitments relating them. The benefit of being able to do that is that logical vocabulary makes it possible to bring the inferential commitments that govern practices of giving and asking for reasons (defending and challenging claims) into those practices as themselves things for which reasons can be given and asked for. Logical vocabulary makes it possible to be critical about the inferential connections between claimables in virtue of which they play the role they do in

---

[5] Of course it always had to be acknowledged that it is not at all clear how this logicist standard for assessing the goodness of theoretical reasoning helps in understanding the distinction between good and bad *practical* reasoning. Rational choice theory has been widely thought to provide a formal analogue on the practical species of reasoning to the role logic plays on for the theoretical species. Latterly, the rapid development of varieties of Bayesianism as a rival framework for assessing the goodness of theoretical reasoning might serve to make our concern with the traditional focus on logic and reasoning seem quaint and archaic. We think logical expressivism provides powerful arguments for philosophers to continue to treat logic as central to their concerns.

reasoning practices, and in that sense mean what they do. Logic should accordingly be understood not as a prescriptive canon for right reasoning, but as an expressive organon: not as providing a standard governing assessments of the correctness of reasoning but as making possible critical investigation and discussion of the credentials of moves as well as positions, inferences as well as claims. Logic should be understood as an organ of critical inferential self-consciousness, and so of critical semantic self-consciousness.

To see how such a view might work, consider two ways of thinking about the correctness of an inference such as:

1) Pittsburgh is to the West of New York.
   So,
   New York is to the East of Pittsburgh.

According to what Wilfrid Sellars calls "the received dogma," one should think of this as an enthymeme: an inference that is only good if one supplies a "missing" premise.[6] Adding the premise:

2) *If* X is to the West of Y, *then* Y is to the East of X.

(and instantiating the variables) turns (1) into an instance of the logically valid scheme of *modus ponens*, detachment from conditionals. That this inference is good is a consequence of the meaning of the concept expressed by the logical vocabulary "if…then__." But another way of thinking about (1) is that it is, in Sellars's terminology, "materially good." By that he means that it is good in virtue of the contents of the *non*logical concepts expressed by the nonlogical vocabulary "West" and "East." Someone who understood those nonlogical concepts, but did not understand conditionals, could still understand that (1) is a good inference. What *makes* (1) a good inference is the content of the nonlogical concepts West and East, *not* any specifically *logical* truth. For those nonlogical conceptual contents are both necessary and sufficient for the inference in (1) to be a good one.

If one thinks about things this way, what should one say about the conditional in (2)? The expressivist answer is that (2) *makes explicit* the goodness of (1)—and of a host of other inferences, such as

3) San Francisco is to the West of New York.
   So,
   New York is to the East of San Francisco.

by codifying a *pattern* of good inferences of which (1) is an instance. It is important to be able to do that. But one need not be able to specify that pattern in order to recognize that (1) is a good inference. One might just have the practical know-how to endorse instances of that pattern, without being able explicitly to *specify* the pattern, to *say* what that pattern is.
On the expressivist view, that is the expressive role characteristic of logical vocabulary. It permits practitioners to codify proprieties of material inferential practice in the form of explicit

[6] Wilfrid Sellars "Inference and Meaning," reprinted in *In the Space of Reasons* [ref.]

principles. To say that the principles are explicit is to say that they are assertible, and so expressible in declarative sentences. In this way they resemble the premises and conclusions of the inferences they license. But they codify, in Mill's terms, "principles in accordance with which to reason" rather than "premises from which to reason." In "What the Tortoise Said to Achilles" Lewis Carroll famously showed the importance of this distinction by showing that conditionals cannot do the work of the rule of *modus ponens*. Sellars urges us to turn the crank on that argument one more time, and treat conditionals as already having the basic expressive role of codifying rules, and only secondarily functioning as premises (and conclusions) of inferences.

According to this order of explanation, the distinction between good and bad reasons is prior to logic. The use of ordinary nonlogical vocabulary already requires distinguishing in practice between claims (or sets of them) that do and do not follow from one another, and those that do and do not rule each other out. Formal logical reason relations presuppose material reason relations.

One can think about the relation between logic and material goodness of inferential relations in two ways: analytically and synthetically. The analytic point of view considers material proprieties of inference in a language that already has both nonlogical and logical vocabulary in it. The inferential relations that hold in virtue of the *logical form* of the sentences involved can then be distinguished from those that hold in virtue of the *material content* of the concepts expressed by the nonlogical vocabulary by the Bolzano-Frege method of observing invariance under substitution.[7] Doing that requires a field of material relations of implication and incompatibility and a distinguished subset of the vocabulary used in the sentences they relate, picked out as the specifically *logical* vocabulary. Then a reason relation can be picked out as holding in virtue of the logical form of the sentences involved just in case two conditions hold:

i. It is materially good.

And

ii. Every relation that results from it by uniformly substituting nonlogical for nonlogical vocabulary is also materially good.

If we pick out the conditional construction, "if…then__" of English as logical vocabulary, then

4) If it is raining, then the streets will be wet.
It is raining.
So,
The streets will be wet.

Holds in virtue of its logical form. For it is a materially good implication, and all substitutional variants of it that keep the conditional construction fixed are *also* materially good implications.

---

[7] In order to get the subjunctive significance right, a more careful statement of the method would quantify over arbitrary extensions of the language that don't alter fundamental grammatical categories.

It is irrelevant to this assessment whether one follows Sellars in taking

5) It is raining.
   So,
   The streets will be wet.

already to be a materially good implication.

Two features of this analytic perspective on the relations between logical and material reason relations are worth remarking. First, it depends on being able to tell logical from nonlogical vocabulary. Quine uses the Bolzano-Frege substitutional strategy to pick out a distinguished class of logical truths from the larger class of truths in general (thereby making himself vulnerable to Dummett's criticism in the passages quoted above). Accordingly, he takes what he calls "the demarcation problem"—what distinguishes specifically logical vocabulary (or the concepts expressed by such vocabulary)—to be one of the principal issues for the philosophy of logic.[8] Expressivism provides a straightforward answer to the demarcation problem: the expressive task distinctive of logical vocabulary is to make reason relations explicit. In the paradigm cases, the defining expressive function of conditionals is to make implication relations explicit and the defining expressive function of negation is to make incompatibility relations explicit.

The second notable feature of the analytic substitutional perspective on the relations between logical and material reason relations is that the methodology of marking invariance under substitution can be applied to *any* distinguished subset of the vocabulary used in the sentences that stand in the material reason relations being considered. One can consider material reason relations that are invariant under substitution of non-culinary for non-culinary vocabulary, or non-nautical for non-nautical vocabulary. Implications and incompatibilities can hold in virtue of their *geo*logical or *astro*logical form. In Dante's theology:

6) Epicurus committed heresy.
   So,
   Epicurus is condemned to the 6th circle of Hell.

holds in virtue of its theological form. This generality of the substitutional analytic methodology raises another challenge for the philosophy of logic: saying why the special subset of reason relations that hold in virtue of their *logical* form is more interesting and important than the subsets of reason relations that hold in virtue of their form with respect to any other sort of vocabulary. This can be thought of as a constraint on responsive answers to the demarcation question. Any such answer must not only pick out the right vocabulary, but must also support an account of its philosophical significance.[9] Expressivism does this by pointing to the critical

[8] Willard van Orman Quine *The Philosophy of Logic* [Harvard University Press, 1970, second edition 1986]. Hilary Putnam agrees in giving the demarcation problem pride of place in his book with nearly the same title, *Philosophy of Logic* [Harper and Row, 1971].

[9] We would argue that prominent views in the philosophy of logic (for instance, the Tarski-Sher approach) falter in the face of this challenge.

rational function of the sort of semantic self-consciousness made possible by the use of vocabulary that plays the expressive role of making reason relations explicit as claims that can themselves be rationally challenged (by offering reasons against them) and rationally defended (by offering reasons for them). The functional role of any bit of vocabulary (whether logical or not) in relations of implication and incompatibility is surely an important aspect of its meaning or conceptual content. (One need not accept the radical semantic inferentialist thesis that conceptual content of an expression just *consists* in the role it plays in such broadly inferential relations in order to acknowledge the semantic significance of those relations.) So demarcating logical vocabulary by its expressive function of making it possible to talk and think critically about those semantically significant reason relations offers a cogent and attractive response to this challenge, too. Adding logical vocabulary to a language that does not contain it brings with it a whole new dimension of critical control, not just over beliefs, but over meanings.

That is the synthetic perspective on the relation between logical vocabulary and reason relations of implication and incompatibility. I call it "synthetic" because it considers the new expressive power that comes from adding logical vocabulary to a prelogical vocabulary-in-use that does not already have the expressive power to codify, and so make it possible to become explicitly aware of and discuss, the reason relations that articulate the meanings of that prelogical vocabulary, rather than analytically picking the logical reason relations out of a larger field that includes both logical and nonlogical implications and incompatibilities. I have characterized the basic thesis of logical expressivism in the philosophy of logic as the claim that the expressive role characteristic of *logical* vocabulary is to make explicit, in the object-language, relations of implication and incompatibility, including the material, prelogical ones that, according to semantic inferentialism, articulate the conceptual contents expressed by nonlogical vocabulary, paradigmatically ordinary empirical descriptive vocabulary. The paradigms of logical vocabulary are the *conditional*, which codifies relations of implication that normatively structure giving reasons *for* claims, and *negation*, which codifies relations of incompatibility that normatively structure giving reasons *against* claims. The synthetic logical expressivist question is: If you don't already have vocabulary performing the expressive role distinctive of logical vocabulary in your language, how can you introduce it?

Further along I'll present our detailed proposal for adding logical expressive power to a set of reason relations defined on a prelogical material base vocabulary. Here I just want to give an initial indication of the ideas that motivate that proposal. Sequent calculus formulations of the rules for logical connectives provide a convenient framework for doing this. To say that a premise-set Γ implies a conclusion A, we can write in the metalanguage: "Γ|~A". (Later on I'll say something about why I am using the funny "snake" turnstile "|~" for implication, instead of the more usual "|—".) As is common in single-succedent sequent calculi, to say that a premise-set Γ is incompatible with a sentence A, we can write in the metalanguage "Γ,A|~⊥", where "⊥" is a symbol for absurdity. (Multisuccedent—and sometimes single succedent—sequent calculi achieve the same effect by using an empty right-hand side.)[10] The expressivist idea is that the expressive task characteristic of conditionals is to make it possible to talk about implications in the logically extended object language: to bring implications into the language as claims that can themselves serve as premises and conclusions of further implications, and so as claims for which reasons can be asked and given.

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

Implication-Codifying Conditional: **$\Gamma |\!\sim A \rightarrow B$ 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.[11]

Such a conditional *says that* an implication holds.

The expressivist idea is that the expressive task characteristic of negation is to make it possible to talk about incompatibilities in the logically extended object language: to bring incompatibilities into the language as claims that can themselves serve as premises and conclusions of further implications, and so as claims for which reasons can be asked and given. So, to perform its expressive task of codifying incompatibility relations in the object language, negation needs to satisfy the

Incompatibility-Codifying Negation: **$\Gamma |\!\sim \neg A$ iff $\Gamma, A |\!\sim \bot$.**

That is, a premise-set implies not-A just in case A is incompatible with that premise-set. This metalanguage statement codifies what is required for the negation of A to express incompatibility with A: That Γ implies the negation of A just in case Γ is incompatible with A. It follows that ¬A is the minimal incompatible of A, in the sense of being implied by everything that is incompatible with A. This is one way to think about the relation between Aristotelian contraries (materially incompatible claims) and Aristotelian contradictories. The contradictory "S is *not* red," formed from "S is red," by applying a logical negation, is implied by all the contraries of "S is red,": "S is green," "S is blue," "S is yellow,"….[12] The logicist tradition, by

---

[10] Such a notational convention conveniently avoids using different symbols for implication and incompatibility, treating all of them notationally as implications. This allows a unified calculus of sequents. Philosophically, it is potentially very misleading. Not only does it invite neglect of the co-equal status of incompatibility with implication as reason relations, but it builds in without comment or justification a crucial structural feature of incompatibility: its symmetry. Later on we will do better and be more careful, by using separate signs for implication and incompatibility. But these introductory remarks do not require breaking with tradition in this regard.

[11] Frank Ramsey "Truth and Probability" (1926), in Frank Ramsey *The Foundations of Mathematics* [Routledge and Kegan Paul, 1931], pp. 158-159.

[12] Strictly, incompatibility is contrariety of *sets* of claims. For instance, in the irreducibly triadic materially incompatible set {"S is a blackberry," "S is ripe," and "S is red"} every pair is incompatible with (contrary to) the remaining sentence.

contrast, treats negation as primitive, and understands material contrariety in terms of formal logical contradictoriness. The contraries of a sentence are all the sentences that imply its contradictory. We see here the opposite explanatory strategies of logicism and expressivism in their starkest form.[13]

The expressive role characteristic of logical vocabulary is visible only from the synthetic point of view, not from the analytic point of view. For that is where the expressive criteria of demarcation presupposed by the analytic methodology of noting invariance under substitution are to be found. That defining expressive role can be further subdivided into an aspect that looks upstream, to the way logical vocabulary is introduced, and an aspect that looks downstream, to the expressive capacity that results from introducing logical vocabulary. These are rough expressive analogues of the circumstances of application and consequences of application Dummettian inferentialism urges us to associate with the use of any vocabulary.[14] For the first, expressivism understands logical vocabulary and the inferential relations that govern its use as *elaborated from* the inferential relations of implication and incompatibility that govern the use of a prelogical material base vocabulary. The logical connective rules used to extend those reason relations from the material base vocabulary to a logically extended vocabulary must require nothing more as input than is present in the use of the base vocabulary. For the second, the logical vocabulary must be introduced in such a way that, when used according to the relations of implication and incompatibility determined by those connective rules, it serves to *explicate* the reason relations that govern *both* the base vocabulary *and* the logically extended vocabulary, in the sense of making it possible to *say* in the extended vocabulary what implies what and what is incompatible with what. In short, logical vocabulary (and its governing reason relations) must be both *elaborated from* and *explicative of* a base vocabulary (and its governing reason relations). For short, we can say it must be "LX" (*el*aborated and *ex*plicative) for the material base.[15]

The explicative expressive function imposes an important criterion of adequacy on the rules for elaborating relations of implication and incompatibility for the logically extended vocabulary from the relations of implication and incompatibility that govern the material base vocabulary. The rules for introducing logical connectives—paradigmatically those that ensure that the conditional codifies implication and negation codifies incompatibility in accordance with

---

[13] Chapter Five of *A Spirit of Trust* [Harvard University Press, 2019] argues that Hegel's concept of "determinate negation" is based on the Aristotelian relation of material incompatibility as contrariety, and that he adopts the order of explanation that proceeds from there to formal logical contradictoriness.
[14] Dummett's idea is further explained in Chapter One of *Articulating Reasons: An Introduction to Inferentialism* [Harvard University Press, 2001], which introduces logical expressivism as a program motivated by semantic inferentialism.
[15] This way of articulating logical expressivism is introduced in Chapter Two of *Between Saying and Doing* [Oxford University Press, 2008].

the principles formulated above[16]—must *conservatively extend* the relations of implication and incompatibility that they elaborate. That is, all the inferential relations among nonlogical vocabulary must be preserved, and no new implications or incompatibilities involving only nonlogical vocabulary must be introduced. Why? Because introducing vocabulary to *express* those reason relations should not *change* them. It should *just* make it possible to express *explicitly* in the (logically extended) object language the reason relations that *implicitly* govern the material base vocabulary (which we theorists express in a proof-theoretic *meta*language of sequents). We will see below that there are *other* notions of explication where it is not appropriate to impose a corresponding conservativeness condition. But not for the one that governs the introduction of specifically *logical* vocabulary. A particularly vivid example of the trouble one can get into if one does not impose this condition is provided by the connective "tonk" introduced by Arthur Prior as part of an argument against natural deduction calculi without restrictions on the rules. He pointed out that the effect of using "tonk" with the introduction rule of classical disjunction, which includes:

$$\frac{p}{p \text{ tonk } q},$$

and the elimination rule of classical conjunction, which includes:

$$\frac{p \text{ tonk } q}{q}$$

then the result licenses the implication:

$$\dfrac{\dfrac{p}{p \text{ tonk } q},}{q}$$

He called that a "runabout inference ticket," since it licenses the implication from arbitrary logically atomic p to arbitrary logically atomic q. Nuel Belnap diagnosed the trouble as a violation of conservativeness.[17] Introducing new logical vocabulary should not create any new implications involving only old vocabulary. Ignoring that constraint courts the risk of "tonking up" one's implication relation.

Although conservativeness is a criterion of adequacy (and so a necessary condition) for any vocabulary playing the distinctive explicative expressive role being recommended as demarcating specifically logical vocabulary, it is not appropriate to require it for nonlogical vocabulary in general. For it is characteristic of ordinary empirical descriptive vocabulary that its use involves endorsing nontrivial material implications relating its circumstances of appropriate application to its appropriate consequences of application. The nonlogical content of

[16] The basic logical system we introduce below, NM-MS, will also include conjunction and disjunction, of course. But from the expressivist perspective, these are essentially auxiliary connectives: Boolean helper-monkeys needed to assist the principal connectives to do their more basic expressive work.
[17] [ref.] Complications to this analysis ensue in substructural settings—in particular if Cut fails.

concepts such as <u>copper</u>, <u>fever</u>, and <u>cruel</u> incorporate many such substantive material implications (and incompatibilities) that need not be redundant relative to other concepts antecedently available. This important dimension of content is ignored and made invisible by understanding conceptual content in terms of truth conditions. For that is the idea of conditions that are both individually necessary and jointly sufficient—the idea of circumstances and consequences of application that are guaranteed to coincide, so that no substantive inferential commitment is involve in the transition between them. The covert ideology behind the notion of truth conditions is accordingly what I have called "logicism." It is the idea that the reason relations that govern *all* vocabulary should not just be *expressible* by logical vocabulary (vocabulary governed by *logical* reason relations), but that those reason relations must be in the end *reducible to* logical reason relations. Since logical reason relations must be conservative over the rest of the language, so must reason relations in general. That is the origin of the requirement of the coincidence of circumstances and consequences of application in the form of conditions that are both necessary and sufficient. But that squeezes out nonlogical content, as articulated by *materially* good relations of implication and incompatibility: those that hold in virtue of the contents of *non*logical concepts. Here logicism about the relations between logic and reason relations misleads us about the latter, and does damage to our semantic theory (according even to a very weak form of semantic inferentialism).

The conditions formulated above for introducing implication-codifying conditionals and an incompatibility-codifying negation give concrete expression to what is required for those logical connectives to count as *explicating* their respective reason relations. Later in the book, Dan Kaplan will show how these core examples—and so the key expressivist idea of the LX-ness of logical vocabulary—can be generalized and made more precise in a rigorous and principled way.

Although being LX for the reason relations of implication or incompatibility that govern some base vocabulary is a *necessary* condition of playing the expressive role characteristic of logical vocabulary, it is not *sufficient*. That is for two reasons. First, the logical vocabulary must be able to make explicit the reason relations that govern the whole logically extended vocabulary, not just the reason relations that govern the material base from which it is elaborated. Second is a generality constraint. To count as genuinely logical, a vocabulary must not just be LX for *some* material base, but for *many*—indeed, for all material reason relations that meet a certain general condition. As to the first, in Frege's first, seminal work, the *Begriffsschrift*, he sets out the task of logic as making conceptual content [begriffliche Inhalt] explicit. And before he made the move Dummett deplores, to focus on truth, he understands such content in terms of inferential role.[18] He aims to be able to articulate the content of nonlogical concepts:

> My concept-script has a more far-reaching aim than Boolean logic, in that it strives to make it possible to present a content when combined with arithmetical and geometrical signs...
>
> Disregarding content, within the domain of pure logic it also, thanks to the notation for generality, commands a somewhat wider domain...
>
> It is in a position to represent the formation of the concepts actually needed in science...[19]
>
> It seems to me to be easier still to extend the domain of this formula language to include geometry. We would only have to add a few signs for the intuitive relations that occur there...The transition to the pure theory of motion and then to mechanics and physics could follow at this point.[20]

He is sufficiently impressed that the logical vocabulary he introduces to codify the contents expressed by nonlogical scientific vocabularies turns out also to be able to specify the inferential roles specified by his own new logical vocabulary that he appends to his book a list showing, for each purely logical proposition he has proven, what other propositions were used in its proof, and what further propositions it is appealed to in proving.

As to the second point, concerning generality, it might be that special disciplines (not just physics and geology, but also astrology and theology) introduce special vocabulary to express inferences unique to them. Such vocabulary, connecting force to mass and acceleration, or blasphemy to damnation, need not for that reason count as logical vocabulary. Logical vocabulary has the expressive task of making explicit the inferential relations of implication and incompatibility that govern *any* autonomous discursive practice—that is, any language-game one could play though one played no other. Neither physics nor theology is an autonomous vocabulary in this sense. The expressive generality logic definitionally aspires to on this expressivist understanding will be important when we discuss substructural varieties of

---

[18] "Right from the start I had in mind the expression of a content...But the content is to be rendered more exactly than is done by verbal language... Speech often only indicates by inessential marks or by imagery what a concept-script should spell out in full."

[Frege, from "Boole's logical Calculus and the Concept-script", *Posthumous Writings* (hereafter *PW*) pp.12-13.]

"...there are two ways in which the content of two judgments may differ; it may, or it may not, be the case that all inferences that can be drawn from the first judgment when combined with certain other ones can always also be drawn from the second when combined with the same other judgments. The two propositions 'the Greeks defeated the Persians at Plataea' and 'the Persians were defeated by the Greeks at Plataea' differ in the former way; even if a slight difference of sense is discernible, the agreement in sense is preponderant. Now I call that part of the content that is the same in both the conceptual content. Only this has significance for our symbolic language [Begriffsschrift]... In my formalized language [BGS]...only that part of judgments which affects the possible inferences is taken into consideration. Whatever is needed for a correct ['richtig', usually misleadingly translated as 'valid'] inference is fully expressed; what is not needed is...not." [*Begriffsschrift* section 3.]

[19] Frege, *PW* p. 46.

[20] Frege, *Begriffsschrift* Preface.

implication and incompatibility, further along. This is not a trivial or non-controversial constraint to impose when demarcating logical vocabulary. Some philosophers of logic restrict its aspirations to codifying the implication and incompatibility relations that govern mathematical proofs, for instance. We side with Frege (and Tarski) here.

When, half a century ago, Quine and Putnam wrote their books on the philosophy of logic referred to above in connection with the demarcation question, they were so comfortably ensconced in logicism about reasoning in general that it did not occur to them so much as to discuss it as a potentially controversial topic. For them the other big issue, besides demarcating logical vocabulary, was the *correctness* question: what is the right logic? The paradigm was the competition between classical and intuitionistic logics, but other candidates included multivalued logics, modal logics, and various other "funny" or otherwise nonstandard logics. What criteria are appropriate to use in choosing between them, and why? And what is the result of such an assessment? Which logic deserves to be certified as the right logic, the one that, as it were, cuts reason at the joints by giving the right answer to the question: what really follows from what?

One notable consequence of offering an expressivist answer to the demarcation question (a question, it was noted above, to which logicism as such has no native candidate answer, though it is compatible with a variety of them) is that the correctness question lapses—or at least, comes to be seen to invite a relaxed, pragmatic, pluralistic response. Consider conditionals. The expressivist idea that conditionals codify implications treats asserting a conditional as endorsing an implication, in the sense of taking it that it is a *good* implication. But there are as many different kinds of conditional as there are dimensions of assessment of the goodness of implications in general. For instance, it is a good thing if an implication does not have true premises and a false conclusion. (At least, it is a bad thing if it does have true premises and a false conclusion.) It is this minimal sense of "good implication" that is expressed by the much-maligned two-valued horseshoe of classical logic. Implications can also be good in the sense that it is *impossible* for the premises to be true and the conclusion to be false. That is the sense of "good implication that is codified by C. I. Lewis's strict-implication hook. Again, it is a good thing about an implication if there is a recipe for turning an argument for the premises (at the limit, a proof of them) into an argument for the conclusion (at the limit, a proof of it). That is the sense of "good implication" codified by intuitionistic conditionals. And so on. These all correspond to different senses of "imply" in the metalanguage. For the expressivist, the question to ask is not which conditional is *correct*—expresses what *really* follows from what. The important question to ask for any candidate conditional is rather exactly what sort of assessment of the goodness of implications it expresses: what sense of "good implication" it codifies. The counsel of wisdom is pluralistic generosity: let a hundred flowers blossom. The richer the logician's armamentarium of expressive tools, the greater the expressive power that can be brought to bear, the greater the number of specialized expressive purposes that can be served.

II. The Structure of Reason Relations: Open or Closed?

I suggested at the outset that the most fundamental question of the philosophy of logic is neither the demarcation question, nor the correctness question, but the question of how logic relates to good reasoning. I have argued that that relation is mediated by reason relations of implication and incompatibility. Those relations provide standards for the normative assessment of reasoning practices, and the expressive task distinctive of logical vocabulary and the concepts it expresses is to make explicit those broadly inferential reason relations. Doing that brings them into our reasoning practices as expressing claims codifying the implicit reason relations, which accordingly can now themselves be critically assessed, defended and challenged by offering reasons for and reasons against them (further claims implying or incompatible with those logically articulated claims).

Once logical vocabulary has been demarcated from other vocabulary in this way, we can pick out some relations of implication and incompatibility that "hold in virtue of the logical form" of sentences formed using logical vocabulary, by the Bolzano-Frege method of noting invariance under substitution. What we get are implications and incompatibilities that are robust under arbitrary substitution of nonlogical for nonlogical vocabulary. These reason relations tell us something important about the content of the logical concepts they articulate—something, we can say, about logical *form*. They don't tell us all about the *content* of logical concepts, because that essentially depends also on the distinctive expressive role those concepts play: the way they serve to make explicit reason relations in general.

Appreciating all of this puts us in a position to ask a further question. What is the structure of reason relations as such? A subsidiary question that then arises is: What is the relation between the structure of *logical* reason relations and the structure of *non*logical reason relations—the implications and incompatibilities that govern the use of ordinary empirical descriptive vocabulary, for instance? Here there is room for a weaker form of what I have called "logicism about reasons." Instead of claiming that "good reason" just means "*logically* good reason," as traditional logicism about reasons does, what we might call "structural logicism" claims that the *structure* of the reason relations of implication and incompatibility that govern *all* reasoning is the same as the structure of *logical* relations of implication and incompatibility. According to this doctrine, although it is conceded that logic does not determine the content that material reason relations confer on nonlogical vocabulary, it *does* determine the structure of implication and incompatibility in general.

What sense of "structure" is in play when we ask this sort of question? It clearly is not the notion of <u>form</u> being addressed when we ask about implications and incompatibilities that obtain in virtue of their *logical* form. For in that sense, the reason relations that structure reasoning with ordinary empirical descriptive vocabulary *have* no form—since they need not involve the use of specifically logical vocabulary at all. The relevant concept of the structure of logical implication and incompatibility was identified relatively late in the development of logic. It was arrived at only in the 1930s, specified by Tarski and Gentzen (in a different form), the founders of the modern model-theoretic and proof-theoretic traditions in logic, respectively. For our purposes here, it will simplify things to start with Tarski's precisification of the algebraic structure of logical implication.[21] His idea was that logical consequence should be understood to have the structure characteristic of topological closure operators. For a countable language L, he considers a consequence relation Cn that assigns to each subset X of the language the set of sentences of L that are its consequences—what, in the notation I have been using are sentences A such that X|~A. He requires first that

1. $X \subseteq Cn(X)$.

That is, all the sentences that are elements of the premise-set count also as consequences of that premise-set. I'll call this principle "Containment" (CO), since it says that the premise-set is contained in the conclusion set.

2. $Cn(X) = Cn(Cn(X))$.

This says that consequence is *transitive*. This condition is the heart of topological *closure*: the consequences of the consequences of any set are already included as consequences of that set. The process of extracting consequences comes to an end, because adding consequences to the premise-set doesn't yield any new consequences.

3. $Y \subseteq X \Rightarrow Cn(Y) \subseteq Cn(X)$.[22]

This says that the consequence operator is *monotone*: adding further premises never take away any consequences. These principles correspond to the Kuratowski axioms for topological closure operators. Containment, transitivity, and monotonicity are also the core structural principles Gentzen imposes on his sequent calculus versions of both classical and intuitionistic logic.[23]

---

[21] I will focus on his initial treatment, in "On Some Fundamental Concepts of Metamathematics" (1930), pp. 30-37 in *Logic, Semantics, and Metamathematics*, Alfred Tarski, J. H. Woodger (trans.) [Oxford University Press, 1956].

[22] Tarski actually uses a different form of this axiom, but immediately proves that it is equivalent to this one.

[23] He calls transitivity "Cut," and monotonicity "Thinning." Instead of containment, he imposes reflexivity—in the notation I have been using, that A|~A. CO follows from reflexivity by monotonicity. Gentzen also imposes three further structural principles: Contraction, Expansion, and Permutation. These serve to turn the lists with which he works into the set with which Tarski works. Since we will stick with premise sets, I will ignore them here, although they are not uncontroversial. In particular, linear logics deny contraction (that if Γ,A,A|~B then Γ,A|~B), with its creator Jean-Yves Girard memorably asking whether twins are two persons or two occurrences of one person, and threatening to give anyone who defends Contraction "two kicks in the ass—not two occurrences of one kick." *Locus Solum*: From the rules of logic to the logic of rules. Mathematical structures in computer science, 11(3):301 (2001).

So by the time Gentzen wrote his classic papers in the mid '30s, there was a consensus, about the algebraic structure of logical relations of consequence or implication. It would prove durable. The trio of Containment (CO), Transitivity (CT), and Monotonicity (MO) was undisputed common ground in debates between proponents of classical logic and intuitionists. One of the reasons that seemed to speak strongly in favor of this characterization is in effect an expressivist one. If the reasoning one wants logic to codify is restricted to mathematical proofs, these structural constraints are just what is needed—whether one understands mathematical proof in classical or in intuitionist terms.

Although I won't pursue the issue in these introductory remarks, which will focus on implication to streamline the discussion, there was a corresponding consensus regarding incompatibility, whose logical form is inconsistency. Both Tarski and Gentzen built into their discussion of consequence a structural principle of *ex falso quodlibet*, sometimes called "explosion." This is the principle that inconsistent premise-sets imply everything in the language. Each exploits that identification to introduce negation—just as the expressivist who understands the essence of negation to be the codification of incompatibility recommends. This move, which was also endorsed both by classical and by intuitionist logicians, had a number of significant consequences. First, it connived, notationally and conceptually, at consigning incompatibility to second-class status among reason relations, relative to implication. Second, it built in without philosophical notice or comment the structural principle that incompatibility relations are *symmetric*: that if a set of sentences (claimables) X is incompatible with a set of setnences Y, then Y is incompatible with X. It did this by treating X and Y as incompatible just in case $X \cup Y$ is incoherent, in the sense that $Cn(X \cup Y)=L$. In this book we will not contest the symmetry of incompatibility. But we do think it is an interesting philosophical question *why* incompatibility is symmetric. We will endorse an argument for and explanation of that claim due to our ROLE colleague Ryan Simonelli. Third, the principle of explosion has been a standing embarrassment both to philosophers of logic, in the Empyrean realm of theory, and to teachers of introductory logic, who have to justify it as "logical" to understandably skeptical students, on the rough ground of pedagogy. We believe it is far too shaky a reed on which to build an understanding of the central reason relation of incompatibility—and so, of the central logical concept of <u>negation</u>. We will instead (in Part Four of the book) ground the account of incompatibility in the normative pragmatics of reasoning practices, in the same terms used to understand implication.

Suppose the traditional consensus about the algebraic structure of logical implication is correct. Relations of logical consequence satisfy containment, transitivity, and monotonicity. (In fact, the logics we will recommend on expressivist grounds do satisfy these global structural conditions.) What about nonlogical or prelogical *material* consequence relations? Do they exhibit the same global algebraic structure that specifically *logical* consequence relations do? The claim that they do, that the global algebraic structure of reason relations generally (not only implication but also incompatibility) is the same as the global algebraic structure of *logical* reason relations might be called "*structural* logicism."[24] Logicism about reasons in general, the view that at least in the theoretical or cognitive arena, *good* reasons just are in the end *logically* good reasons, would seem to entail structural logicism. If and insofar as that is so, any argument against structural logicism about reason relations is an argument against logicism about reasons *tout court*. Perhaps there is enough daylight between these positions that one could consistently maintain logicism about reasons and reason relations while accepting that reason relations in general do not satisfy the same global structural constraints that *logical* implication and incompatibility (inconsistency) do. But it is hard to see how the motivations for endorsing logicism about reasons generally would survive such a concession and the heroic measures needed to reconcile it with acknowledging the substructurality of material reason relations.

In any case, there are good reasons to think that implication in general does not have the structure of a topological closure operation that Tarski found to characterize the specifically *logical* implications on display in mathematical proofs. Of the three structural principles that articulate the canonical consequence-as-closure conception, the one that has aroused the most skepticism is monotonicity: the principle that the implication of a conclusion by some premise-set is never infirmed or defeated by the addition of further premises. Containment, the principle that among the consequences of any set of premises are to be found those premises themselves, has seemed at best obviously correct and at worst a harmless concession, stipulation, or *façon de parler*. Transitivity, the principle that implications can be strung together, with the conclusions of some serving as the premises of others, has seemed to be presupposed by the possibility of extended consecutive reasoning. (While it would be implausible as a descriptive claim and disastrous as a prescription to rule out reasoning like that, failing to impose it as a *global* constraint would not require denying that *sometimes*, locally, implication *is* transitive. To say that sometimes implications cannot be strung together does not entail that they never can. It just raises the question of what distinguishes the two cases.) While both global structural principles have been questioned (and we will join in on some of the complaints)—and not only as manifestations of the pathological suspiciousness, skepticism, and paranoia of those who have spent too much time wrestling with semantic paradoxes—it is safe to say that the structural principle whose applicability to ordinary reasoning has come in for the most criticism is monotonicity.

Monotonicity is just not a plausible constraint on *material* consequence relations in general. Outside of mathematics and perhaps fundamental physics, almost all actual reasoning is defeasible. Usually when good reasons are offered supporting a conclusion, the acquisition of

---

[24] Tarski at least gives aid and comfort to this view by sometimes dropping the qualification "logical" that defines his official topic, and just talking about "consequence" relations,, in his 1936 paper on "On the Concept of Logical Consequence,"[24] which offers a semantic, model-theoretic account of the notion of (at least) logical consequence, aiming to show *why* it has the algebraic properties it does.

further information can undercut that support. This is true in everyday reasoning by auto mechanics and on computer help lines, and in more institutionalized, higher-stake forms of reasoning conducted in courts of law and in medical diagnosis. Rare is the argument in these contexts whose conclusion cannot be contested without contesting any of its premises, by appending an "unless…" clause. The same holds for probabilistic reasoning, in which additional information serves to alter the reference class with respect to which frequencies are assessed. New information can shrink the reference class, expand it, or just shift it, and such changes can both lead to new conclusions and infirm old ones. If all that is known of some particular organism is that it was chosen at random from all the organisms on earth, then it is probably a bacterium, because they make up such a high proportion of those organisms. If the information is added that the organism is multicellular, then it is surely not a bacterium, but is probably a plant. Add the data that it is free-moving, then it is probably marine. Add that it is terrestrial, then it is probably an insect. If in addition it is larger than a breadbox, then it is probably a vertebrate. And so on.

Since probabilistic reasoning is broadly inductive, it might seem that one could hold onto monotonicity by restricting logic to deductive implication relations. If one builds monotonicity into the idea of deduction, by requiring derivation according to rules defining a closed consequence relation, this would be straightforwardly circular. If not, then the restriction to deductive implication relations must mean something like "dispositive" as opposed to "probative" reason relations—that is, those that govern committive, rather than merely permissive reasoning, reasoning where the premises necessitate the conclusion, rather than merely rendering it likely. But defeasibility, hence nonmonotonicity of implication relations, is a structural feature not just of probative or permissive reasoning, but also of dispositive, committive reasoning. For instance, nonmonotonicity is a well-known feature of subjunctive reasoning. Indeed, one of the substantial criteria of adequacy for semantic theories of subjunctive conditionals codifying such reasoning is that they can make sense of "Sobel sequences," in which further information sequentially flips the valence of the implication. If were to I strike this dry, well-made match, it would light. But *not* if it is in a very strong magnetic field. Unless, additionally, it were in a Faraday cage, in which case it would light. But *not* if the room were evacuated of oxygen. And so on. Notice here that reasons *against* a claim are as defeasible in principle as reasons *for* a claim. Material incompatibility relations are no more monotonic in general than material implication relations. Examples of the one sort readily convert into examples of the other sort. Claims that are incompatible in the presence of one set of auxiliary hypotheses can in some cases be reconciled by suitable additions of collateral premises. Cases with this shape are not hard to find in the history of science.

Why is the structure of material reason relations like this? Nonmonotonicity is grounded in the demands of ordinary reasoning practices. In actual practice, interlocutors must be able to state their reasons for accepting or rejecting conclusions. The implication relations that govern such practices accordingly have finite premise-sets. One cannot enumerate all the possible conditions under which the implication would not hold. And the problem is not even exactly that the additional premises that would infirm the implication form an infinite set. Worse, the possible defeaters—the considerations that would need to be enumerated in “unless” clauses to make the implication water-tight and indefeasible—are typically not even a definitely specifiable set. Rather, it is indefinitely extensible, always open to the discovery of other ways the implication could go wrong. (The nutritious food won’t be edible if it is microscopic or gigantic, encased in glass, has a finkish disposition to vanish or turn to poison or set off an explosion if touched….) I’ll return to the issue of the attempt to make the invocation of an infinite set (of possible defeaters or consequences) take the place of what is really an open-ended process further along, in discussing further what is wrong with the idea of consequence relations as closure operators.

It is sometimes thought, or at least hoped, that the problem of almost every empirical implication having a class of defeaters that defies specification can be solved by quantifying over them using *ceteris paribus* (“all things being equal”) clauses. This strategy is a mistake, and rests on a confusion. The result of appending such a clause to an implication would be to trivialize the claim if it has the effect of saying “Premise-set S implies conclusion A—unless, for some reason, it does not.” Nor does the idea that the addition of a *ceteris paribus* clause suffices to turn nonmonotonic implications into monotonic ones make sense. The proper term for a Latin phrase whose recitation can do *that* is “magic spell.” The fantasy of changing the algebraic structure of implication by waving a wand is the result of radically misunderstanding the expressive function of *ceteris paribus* clauses, which is explicitly to *mark* and *acknowledge* the defeasibility, hence nonmonotonicity, of an implication, not to cure it by *fiat*. One appends such a phrase to the reasons one offers, the implication one endorses, or the conditional one asserts just to admit that one knows there are further conditions, which if they obtained would defeat the implication—perhaps with the pragmatic implicature that one does not know of any that actually do hold.

Suppose that, in view of all these features of ordinary, nonlogical reasoning, it is agreed that that among the senses of “implication” that expressivists should aspire to codify are nonmonotonic ones. Then relations of consequence (following from or being a reason for) in general do not have the full structure of closure operations—as the tradition of Tarski and Gentzen takes it that specifically *logical* consequence relations do. That is to give up *structural* logicism, as far as monotonicity is concerned. Is there a weaker structural condition in the vicinity that can be put in its place? Containment (CO) says that in the very special case of consequences of implications that are also premises of those implications, one can weaken the premise-set with arbitrary additional premises without infirming the conclusion. But that is a very special case, amounting to a kind of triviality. Denying monotonicity (MO) means not assuming that one can weaken *every* implication with arbitrary additional premises without infirming the conclusion. Is there a way to specify some more restricted set of additional premises one could add that would be guaranteed not to defeat an implication? One plausible candidate answer to this question is given by what has been called "*cautious* monotonicity." Cautious monotonicity says that while one might indeed not be able to weaken an implication by adding just any sentence as a collateral premise without running the risk of infirming it, it should at least be safe to add further premises that are already implied by the original premise set.

Cautious Monotonicity (CM):

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

The idea is that since the premise set Γ already implies A, adding A an explicit premise should not cause any trouble with other consequences of Γ. Even though there might be *some* additional premises that *would* infirm the implication, sentences that are *already implied* by the premise-set are not among them.

It has often been argued not only that cautious monotonicity is a plausible principle, but that it is in effect indispensable: that it is a *minimal* condition that well-behaved nonmonotonic consequence relations must satisfy.[25] Satisfying CM is generally regarded as a criterion of adequacy for assessing nonmonotonic logics. CM plays a prominent role, for instance in what Kraus, Lehman, and Magidor call the "core properties" or the "conservative core" of non-monotonic systems (and for this reason are now often called the "KLM structural properties" required of nonomonotonic systems), and count it a signal virtue of their preferential semantics for nonmonotonic logic that it validates this structural principle.[26] By contrast to monotonicity, there has not been much skeptical philosophical attention directed at cautious monotonicity. This is a shame, because the underlying issues here are just as important, and addressing them is deeply revealing of considerations that remain invisible if the discussion remains at the level of the much stronger structural principle of monotonicity.

The first step in appreciating this is realizing that cautious monotonicity is the dual of cumulative transitivity, a version of Gentzen's "Cut."[27] This structural principle is expressed in Tarski's algebraic metalanguage for consequence relations by the requirement that the consequences of the consequences of a premise-set are just the consequences of that premise-set, and by Gentzen as the principle that adding to the explicit premises of a premise-set something that is already part of its implicit content does not add to what is implied by that premise-set. It is the principle appealed to in chaining together implications in extended consecutive reasoning.

[25] This case has been made forcefully by Dov Gabbay [Gabbay, D. M., 1985, "Theoretical foundations for nonmonotonic reasoning in expert systems", in K. Apt (ed.), *Logics and Models of Concurrent Systems*, Berlin and New York: Springer Verlag, pp. 439–459.], who includes also CO and CT as necessary for workable nonmonotonic systems.

[26] Kraus, Sarit, Lehmann, Daniel, & Magidor, Menachem, 1990. Nonmonotonic Reasoning, Preferential Models and Cumulative Logics. *Artifical Intelligence*, 44: 167–207.

[27] "A version of" because CT is additive, that is, context sharing, while Gentzen's Cut was multiplicative, that is, context combining. In open, substructural settings, these diverge.

Cumulative Transitivity (CT): $\dfrac{\Gamma \mid\!\sim A \qquad \Gamma, A \mid\!\sim B}{\Gamma \mid\!\sim B.}$

CT says that adding a consequence of a premise-set to that premise-set never *adds* consequences—that what a premise-set implies when we add its own consequences to it already implies all on its own—while CM says that adding a consequence to the premise-set never *subtracts* consequences the original premise-set had.

Here is a way to think about the underlying issue. Using language that was second-nature to Leibniz and Kant, we can think about the *content* of a set of claimables in the literal sense of what is *contained in* it. A set $\Gamma=\{A_1, A_2,\ldots A_n\}$ literally contains all of the sentences A_i in the set-theoretic sense that these are the elements of the set Γ. We may say that it contains them *explicitly*, since they are what we specify when we specify the set. They are the *explicit content* of the set. If it now happens that Γ implies A—in our notation, Γ |~ A—then we can say that A is *implicit in* Γ, in the literal sense of being *implied by* it. A, then, is part of the *implicit content* of Γ. (Analogously, we might think of every set Δ that is materially *incompatible* with Γ as being part of Γ's *contrastive* content.) Then CM and CT can be thought of as having a common topic. Both concern what happens when the status of some consequence of a premise-set is changed, by turning it into an additional premise. The process of moving a sentence from the right-hand side of the implication turnstile to the left-hand side, from appearing as a conclusion to appearing as a premise, might be called the process of *explicitation*. For it is the process of making some *implicit* bit of content *explicit*, turning what is *implicitly* contained in a premise-set into something that is *explicitly* contained in it. Explicitation in this sense is not at all a *psychological* matter. And it is not even yet a strictly *logical* notion. For even *before* logical vocabulary has been introduced, we can make sense of explicitation in terms of the structure of *material* consequence relations. Noting the effects on implicit content of adding as an explicit premise sentences that were already implied is already a process available for investigation at the semantic level of the *prelogic*.

Both CT and CM concern the effects that explicitation has on the *consequences* of the premise-set, comparing the consequences before explicitation with the consequences after explicitation. Since the consequences of a premise-set are its implicit content, CT says that explicitation does not *gain* any implicit content, and CM says that explicitation does not *lose* any implicit content. CT says no consequences are added, and CM says not consequences are subtracted. Together, they entail that ***explicitation is inconsequential***: making implicit content of a premise-set explicit has no effect on its consequences at all. Moving a sentence from the right-hand side of the implication-turnstile to the left-hand side does not change the consequences of the premise-set. It has no effect whatever on the implicit content, on what is implied.

I began by opposing an *expressivist* approach to understanding the relations between logic and reasoning (mediated by reason relations of implication and incompatibility) to a *logicist* approach to understanding them, according to which all good reasons are *logically* good reasons: every genuine implication is valid in virtue of the logical form of its premises and conclusion. I then considered a weaker, purely *structural* form of logicism. It claims that the algebraic structure of material reason relations of implication and incompatibility is the same as the algebraic structure of specifically *logical* relations of implication and incompatibility. For historical reasons I have gestured at, philosophers of logic have taken that algebraic structure to be topological closure. Topological closure is a matter of satisfying monotonicity and transitivity, MO and CT (as well as containment, CO, or at least reflexivity, RE). I am now claiming that it is more philosophically revealing to focus on a different kind of closure structure, which involves pairing CT not with MO, but with CM. This might be called "*explicitation closure*," since it entails that explicitation is inconsequential. Since MO entails CM, rejecting the explicitation-closure form of structural logicism—by denying that explicitation is inconsequential for material consequence relations—will entail rejecting the topological-closure form of structural logicism. That is the view I want to argue for now.

It might well be sensible to require the inconsequentiality of explicitation as a structural constraint on *logical* consequence relations. But just as for the logical expressivist there is no good reason to restrict the rational relations of implication and incompatibility we seek to express with logical vocabulary to monotonic ones, there is no good reason to restrict our expressive ambitions to consequence relations for which explicitation is inconsequential. On the contrary, there is every reason to want to use the expressive tools of logical vocabulary to investigate cases where explicitation *does* make a difference to what is implied.

One such case of general interest is where the explicit contents of a premise-set are the records in a *database*, whose implicit contents consist of whatever consequences can be extracted from those records by applying an *inference engine* to them. (The fact that the "sentences" in the database whose material consequences are extracted by the inference engine are construed to begin with as *logically* atomic does not preclude the records having the "internal" structure of the arbitrarily complex datatypes manipulated by any object-oriented programming language.) It is by no means obvious that one is obliged to treat the results of applying the inference-engine as having exactly the same epistemic status as actual entries in the database. A related case is where the elements of the premise-sets consist of experimental *data*, perhaps measurements, or observations, whose implicit content consists of the consequences that can be extracted from them by applying a *theory*. In such a case, explicitation is far from inconsequential. On the contrary, when the CERN supercollider produces observational measurements that confirm what hitherto had been purely theoretical predictions extracted from previous data, the transformation of rational status from *mere* prediction *implicit* in prior data to actual empirical observation is an event of the first significance—no less important than the observation of something incompatible with the predictions extracted by theory from prior data. This is the very nature of empirical *confirmation* of theories.

Imposing Cut and Cautious Monotonicity as global structural constraints on material consequence relations amounts to equating the epistemic status of premises and conclusions of good implications. But in many cases, we want to acknowledge a distinction, assigning a lesser status to the products of risky, defeasible inference. In an ideal case, perhaps this distinction shrinks to nothing. But we also want to be able to reason in situations where it is important to keep track of the difference in status between what we take ourselves to know and the shakier products of our theoretical reasoning from those premises.

***

Let us take stock. In the first section of this Introduction I considered two diametrically opposed approaches to understanding the relations between logic and reasoning. Taking on board the idea that logic concerns the reason relations of implication and incompatibility that govern reasoning practices and processes sharpened the issue somewhat. Logicism claims that logic determines the proper relations of implication and incompatibility: that implications hold just in case they are or are supported by *logically* good implications, and that incompatibilities are or are supported by *logical* incompatibilities, that is, formal inconsistencies. Expressivism understands the task of logic to be expressing material, prelogical reason relations of implication and incompatibility: making them explicit in the sense of sayable, claimable contents, for which reasons can be asked and given.

I began the second section by considering a weaker form of logicism: *structural* logicism. This is the view that the algebraic structure of material reason relations is the same as the algebraic structure of specifically logical reason relations: implications and incompatibilities that hold just in virtue of the logical form of the sentences involved. For the case of logical consequence, it is generally agreed that it has the structure of a topological closure operation. Combining Tarski's and Gentzen's versions of these structural principles, so as to extrude irrelevant details particular to their formulations, this means that consequence relations satisfy Containment (CO), Monotonicity (MO), and Cumulative Transitivity (CT). The traditional form of structural logicism accordingly takes implication to have a *topological closure* structure. Assuming Tarski and Gentzen are right about the structure of logical consequence (and this much is not at issue, for instance, between classicists and intuitionists), logicism about implication generally entails the topological closure form of structural logicism. If that is wrong about implication in general, if material consequence relations do not in general exhibit the topological closure that comprises CO, MO, and CT, then logicism cannot be right either.

I then argued that, however it might be with specifically *logical* implication, *material* consequence relations are not in general monotonic. That is enough to show that the topological closure version of structural logicism is not true. But what structure do consequence relations in general exhibit? It was pointed out that a popular principled fallback from MO is Cautious Monotonicity (CM).[28] It is the principle that no consequences of a premise set Γ are lost by adding any collateral premises that are already consequences of Γ. CM seems a particularly natural weakening of MO to consider, because it is dual to CT. CT says that adding consequences of Γ to Γ never *adds* any new consequences, and CM says that adding consequences of Γ to Γ never *subtracts* any consequences. Observing this duality brings into view the operation they share: moving a sentence across the implication turnstile, from being a conclusion to being a premise. Thinking of what a premise set implies as what it contains *implicitly* (its implicit content) and the actual elements of the premise set as what it contains *explicitly* (its explicit content) makes it natural to call this process "explicitation." It is making implicit content explicit: a prelogical sort of expression.

Explicitation, in turn, makes visible a further kind of closure structure: explicitation closure. If both CM and CT hold globally for an implication relation, then explicitation is guaranteed to be *inconsequential*. Making implicit content explicit never affects implicit content, neither increasing or decreasing it. This sort of closure is weaker than topological closure, just insofar as CM is weaker than MO. But it, too, expresses commitment to a kind of *stability* of consequences, a kind of closure.

And with this weaker sort of closure structure comes a new sort of structural logicism: explicitation closure structural logicism. This is the claim that consequence relations in general have at least this much of the structure of logical consequence: CM and CT, as well as CO.[29] Explicitation closure might indeed seem to be an attractive fallback from topological closure as a candidate for being the structure of consequence relations generally. It is implied by the full topological closure Tarski and Gentzen take to be characteristic of logical consequence, and

[28] I do not discuss the other most popular candidate weaker than MO, often called "rational monotony." because as usually formulated, it assumes the language already has negation in it. It says that no conclusions of a premise set Γ are lost by adding new premises that do not contradict Γ. There is a version that appeals only to incompatibility, but looking at structural principles relating implication and incompatibility would take us too far beyond the argument of this Introduction.

[29] CO comes from Tarski's plausible condition that $X \subseteq Cn(X)$: any premise set is contained in its consequence set. In the context of MO, CO is equivalent to Reflexivity (RE), which is what Gentzen actually uses (all leaves of all purely logical sequent derivations are instances of RE). If we relax MO, this equivalence breaks down, and one might worry, as relevance logicians do, about endorsing even the weak sort of monotonicity that CO enforces: monotonicity of implications of the form A|~A, which can be arbitrarily weakened with further premises. But CO remains plausible when thought of in explicitation terms: what a premise set *explicitly* contains counts as also *implicitly* contained in it. Explicit content is part of implicit content.

would permit acknowledgment of a distinction from material consequence relations, if they only satisfy the weaker closure structure.

Note that MO is as well-defined for incompatibility as for implication. An incompatibility property is nonmonotonic if adding elements can make incompatible sets compatible. And monotonicity is equally implausible as a constraint on material reason relations. Indeed, there is a general procedure for turning failures of MO for material consequence relations into failure of MO for material incompatibility relations. Consider the Sobel sequence gestured at as an example of nonmonotonicity above. Striking the dry, well-made match and its lighting is incompatible with its being in a strong magnetic field. But striking the dry, well-made match and its lighting and its being in a Faraday cage *is* compatible with its being in a strong magnetic field. Negation introduced to codify material incompatibility relations will have to deal with the structural nonmonotonicity of incompatibility every bit as much as conditionals introduced to codify material consequence relations will have to deal with the structural nonmonotonicity of implication. That this observation is of some importance becomes clear in light of the fact that prominent approaches to nonmonotonic logics (such as preferential models and default reasoning) help themselves to classical negation and specifically, the monotonic relation of inconsistency it supports.[30]

But there is no analogue of explicitation for incompatibility. Nothing stands to incompatibility as explicitation stands to implication. We'll see in Part Four of the book that there are different structural issues unique and native to incompatibility: namely, symmetry.

To argue against the explicitation closure form of structural logicism I introduced a special case of consequential reason relations to serve as a model. In Part Four of the book, we will offer some suggestions as to what "follows from" (and "incompatible with") should be understood to mean in general. But for present purposes it is helpful to think of the explicit premises of an implication as a database, and the turnstile as standing for a theory functioning as an inference-engine that, when applied to the database, yields the conclusions, thereby extracting the content implicit in the database. Thinking instances where the database contains observational data and the inference-engine extracts the predictions of some theory shows that

---

[30] Kraus, Sarit, Lehmann, Daniel, & Magidor, Menachem, 1990: Nonmonotonic Reasoning, Preferential Models and Cumulative Logics. *Artifical Intelligence*, 44: 167–207, John F. Horty, 2007: Defaults with Priorities. *Journal of Philosophical Logic*, 36: 367–413.
As explained below, their projects of building nonmonotonic logics out of classical monotonic ones is alien to the expressivist approach. We want logics that are expressively adequate to codify nonmonotonic reason relations of implication and incompatibility. That is not at all the same enterprise. In fact it turns out that the purely logical reason relations governing such logics (the implications and incompatibilities that hold in virtue of logic alone) can be structurally closed—indeed monotonic and so topologically closed. So in the end we are in a position to justify the invocation of classical negation and inconsistency as part of the definition of one's nonmonotonic logic—if one still wants to do what these nonmonotonic logicians want to do.

the inconsequentiality of explicitation is *not* a condition we want to insist on in general. The difference in status between what has been observed or measured and what is merely predicted by theory is too important to have the boundary between them structurally erased. So this weaker form of closure should also be rejected as a global structural constraint on consequence relations, and explicitation closure structural logicism must accordingly be rejected along with topological closure structural logicism.

It will have been noticed that the arguments against the global inconsequentiality of explicitation for consequence relations in general focused on CT rather than CM—that is, on the idea that making inferentially implicit content explicit could never make it possible to derive consequences that could not be derived before that explicitation. (It might be worth pointing out that the most common philosophical motivation for denying Cut arises from consideration of the semantic paradoxes, and the argument here depends on no such "funny business.") But what about CM? Can it also happen that confirming *some* conclusions extracted by theory from the data infirms *other* conclusions that one otherwise would have drawn? I think a case could be made out for the intelligibility and coherence of inference-engines that allow this. But arguing it is not to my purpose. For we must not lose sight of the issue that led to addressing the issue of whether material consequence relations should be understood to be structurally closed in the first place (whether in the topological or the explicitation sense). It was to understand the constraints on *expressing* material reason relations using *logical* vocabulary. The database + inference engine model reminds us that we should aim for the greatest possible flexibility and expressive power possible. We should build as few *a priori* constraints on logically codifiable inference engines as possible. The expressivist's aim should be to produce and deploy logical tools for expressing reason relations of all intelligible structures. The expressivist ideal is to develop the expressive power to make explicit *any* and *all* species of the turnstile, any and all senses of "follows from." Thinking in terms of databases and inference engines reminds us of just how capacious that class (and so that aspiration) is. Perhaps it is a utopian aspiration. (Spoiler: It is not.) From *this* point of view, CM should not be assumed to hold globally, any more than CT should.

It is widely recognized that failures of monotonicity generate failures of simple transitivity. So, a standard example of an MO failure is:

Tweety is a bird, so (probably) Tweety can fly.

But *not*

*Tweety is a bird and Tweety is a penguin, so (probably) Tweety can fly.

The corresponding failure of simple transitivity is

Tweety is a penguin, so Tweety is a bird.
Tweety is a bird, so (probably) Tweety can fly.

But *not*

Tweety is a penguin, so (probably) Tweety can fly.

Ryan Simonelli of our ROLE group has observed further that there is a general procedure for turning examples of failures of monotonicity into examples of failures of *cumulative* transitivity (CT): cases where explicitating a consequence adds further consequences.
It is *not* the case that

*Tweety is a bird, so Tweety is a non-penguin.

But we do have

Tweety is a bird, so (probably) Tweety can fly.

Tweety is a bird and Tweety can fly, so Tweety is a non-penguin.

Adding the consequence (fly) to the premise set (bird) that does not imply non-penguin yields a premise set that *does* imply non-penguin.
In this way we cans see that where one finds nonmonotonicity in material consequence relations, one will also find failures of CT. (Note that this sort of argument, too, does not involve appeal to semantic paradoxes.)

In fact, expressivism offers an even better argument against explicitation closure (and so, this weaker form of structural logicism) than we get even from the database + inference engine model. This is the observation, due to Ulf Hlobil, that CT (in the context of CO), together with an implication-codifying (double) Ramsey conditional forces MO. [31] For if we start with some arbitrary implication Γ|~A, we can derive Γ,B|~A for arbitrary B—that is, we can show that arbitrary additions to the premise-set, arbitrary weakenings of the implication, preserves those implications. And that is just monotonicity. For we can argue:

Γ|~A	Assumption
Γ,A, B|~A	CO
Γ,A|~B→A	Ramsey Condition Right-to-Left
Γ|~B→A	CT, Cutting A using Assumption
Γ,B|~A	Ramsey Condition Left-to-Right.

As a result, a proper conditional cannot be introduced conservatively on a nonmonotonic base. Such a conditional cannot in principle *explicate* (conservatively express) a nonmonotonic base. If we want such a conditional (and CO), we must forego CT as a global principle. For it is not just that CT creates a problem (forces monotonicity) if the language already contains an implication-codifying conditional. The problem is that we cannot *add* such a conditional to a nonmonotonic base language without endorsing new (monotonic) implications involving only the old, prelogical vocabulary. And that violates the explicative, implication-codifying expressive task characteristic of conditionals. From an expressivist point of view, this is decisive: we need a logic that can be introduced conservatively over, and has the expressive power to codify, material reason relations that are *open*, not just in not being *topologically*

[31] Hlobil, U. (2016), "A Nonmonotonic Sequent Calculus for Inferentialist Expressivists." In Pavel Arazim and Michal Dančák (eds.) The *Logica* Yearbook 2015, pp. 87-105, College Publications: London.

closed, but in not being *explicitation* closed, either. Structural logicism in both forms must be rejected.

The material consequence relations the expressivist takes it logic should aim to codify are accordingly radically substructural, in that it should not be presupposed that they satisfy global structural principles of the sort characteristic of specifically *logical* consequence relations. The kind of structure denied is *closure* structure, of the two sorts distinguished here: topological closure and explicitation closure. The substructural consequence relations they contrast with are *open*. The significance of reason relations with open structure is best grasped by thinking about the process of explicitation. Explicitation as a process is *drawing* or extracting consequences, and adding them to the premises from which one reasons. (Harman reminds us that this is not the only rational process governed by implication relations.) The key point is that in an open structural setting, making explicit any set of consequences of a premise set might add some new consequences (where there are local failures of CT) and subtract others (where there are local failures of CM), relative to the consequences of the original premise set. This has a number of consequences, which highlight the striking differences with structurally closed consequence relations.

First of all, if closure is not required, there is no guarantee that the process of explicitating consequences, and then explicitating the consequences of *that* expanded premise set, and then explicitating the consequences of *that* set will reach a fixed conclusion. In open settings one cannot be sure in advance that the process of explicitation will reach a stable stopping-place—that it will arrive at a premise-set all of whose explicitations are inconsequential, involving no violations of CT or CM. It can happen that every position $X \cup Y$ one arrives at by explicitating consequences of the original premise set X (CO guarantees one will only arrive at supersets of X) still has some implicit content, some set of consequences, such that when they are added to $X \cup Y$ as explicit premises, results in different consequences than $X \cup Y$ had. The process of explicitation need not end. The explicitation closure conditions, by contrast, guarantee *finality*: that Cn(X)=Cn(Cn(X) for every premise set.

Secondly, in particular, with open consequence relations there is nothing privileged about the result of explicitating *all* of X's consequences at once. One might get a larger implicit content by explicitating only some of the consequences of a premise set. Closed consequence relations ensure that the result of explicitating all the consequences of a premise-set will include the result of explicitaing only some of them. So not only is finality guaranteed, but one can get to the final fixed point of the explicitation process in a single step. Explicitation is *immediate*.

Thirdly, explicitation of consequence relations with open structure can be radically path-dependent. Nor is there any guarantee that on an explicitation path that starts by adding those consequences one will ever arrive at any of the premise sets or consequence sets reached by starting with the explicitation of a different proper subset of the consequences of Γ. As a result, the consequences that come into view at any point in the process of explicitating some premise set Γ depend on which of Γ's consequences one chooses to explicitate first. This is explicative *hysteresis*. By contrast, explicitation of structurally closed consequence relations is *stable*, in the sense of being path-independent. Since it is *final* there is a fixed endpoint to the process: Cn(X). Since it is *immediate* one can jump to that endpoint in a single explicitating step. And since it is *stable*, if one *did* explicitate step-wise, the results would be the same endpoint, for all explicitation paths from X to Cn(X).

Finality, immediacy, and stability of explicitation are all very useful features of consequence relations. There are many good reasons to want to build them into the *logical* consequence relations that govern the use of the logical vocabulary whose expressive task it is to make explicit the reason relations that govern the use of all kinds of vocabulary. But if, through a thoughtless and misplaced commitment to structural logicism, we project those ideals onto the actual material consequence relations that govern the use of nonlogical vocabulary and articulate the conceptual contents they express, we make invisible the crucial rational *process* of *drawing* consequences from premise sets, of acknowledging explicitly the implicit, consequential contents of explicit commitments. The process of explicitation is important. Studying it is one of the reasons we want logical tools to make reason relations explicit. The most elementary sort of pragmatism counsels that we not obscure rational processes and practices by imposing Procrustean *a priori* structural restrictions.

From the expressivist point of view, it is important to keep in mind that there are really *three* levels of consequence relation that must be clearly distinguished, and which have different algebraic structures. At the ground level are material consequence relations. These are structurally open, nonmonotonic, and not globally transitive. (Material incompatibility is also not in general globally monotonic.) On the basis of those material reason relations, specifically *logical* vocabulary is introduced, to make explicit those relations of implication and incompatibility. To perform its distinctive expressive function, the reason relations governing logical vocabulary must be elaborated from the material reason relations *conservatively*—that is, so that no new implications or incompatibilities involving only nonlogical vocabulary are forced by the rules that introduce logical connectives. Since the underlying material base implications and incompatibilities are structurally open, so must the overall reason relations governing the logically extended base language be. However, that is compatible with full structural closure at the third level. The purely logical consequences and incompatibilities, those that hold in virtue of their logical form alone, can be fully monotonic and transitive. It is perhaps surprising that full closure structurality of the logical fragment of the reason relations governing the logically extended base language is compatible with the logical vocabulary performing its essential expressive function of making explicit *open*, *sub*structural reason relations in the base vocabulary. But, as we shall see in Part Two of the book, it is so. (By contrast to more traditional approaches, we offer not a nonmonotonic logic—in a straightforward sense, the logic we recommend is fully monotonic and transitive—but a logic of nonmonotonic and nontransitive reason relations.) Structural logicism, whether of the stronger topological closure variety or of the weaker explicitation closure sort, is results from not clearly distinguishing the different consequence relations in this heirarchy, and not appreciating the different structural demands that are appropriate for each.

***

There are three large questions about the macrostructure of reason-relations that remain after we have rejected structural logicism, by refusing the Procrustean projection either of topological closure or of explicitation closure onto reason relations in general, just because they hold of specifically *logical* relations of consequence and incompatibility. The first is really a metastructural question: Why are there *two* reason relations, implication and incompatibility? Why not just one, or three? And if there must be two, why just *these* two? This is an important, indeed fundamental question, even though it has not attracted the philosophical interest and attention it deserves. Tarski only addresses the structure of logical consequence relations, and Gentzen the structure of sequents. We opened this Introduction by quoting Dummett identifying the subject-matter of logic as "the relation of logical consequence." But logical inconsistency is as fundamental to logic as logical consequence—and that is true independently of whether one is a logicist or an expressivist about the relations between formal logical and material reason relations. Perhaps the neglect of this issue by philosophers of logic is the result of its seeming obvious because baked into a bivalent approach to reasoning that begins with the opposition between truth and falsity. Certainly it was encouraged by the habit of encoding inconsistency into the consequence relation, either by explosion (*ex falso quodlibet*) or by sequents with empty right-hand sides. We would argue that this notational convenience should not be allowed to tempt one into thinking that implicational explosion is of the essence of incompatibility generally, and that if it is not, the question of how to understand the relations between incompatibility and implication becomes visible as substantial and even urgent.

The other two structural questions that then arise within the scope of this macrostructural question then concern the two reason relations. If implication and incompatibility in general—that is, material consequence and material incompatibility, as opposed to their refined formal logical species—do not exhibit the closed algebraic structures traditionally attributed to them, is there anything general we can say about the open structures they *do* exhibit? For implication, what is left after we have rejected the imposition (the importation from logical consequence) of global transitivity and monotonicity—and even the cautious monotonicity that is dual to cumulative transitivity? For material incompatibility, we reject monotonicity as well. There is no analogue of transitivity here, but there is not because it appears to be *de jure* a *symmetric* relation, just as logical notations build in. Is this *just* logicism about incompatibility? Or is it the case that material incompatibility is in general symmetric? Note that this is in contrast to implication, which in general is not. Must material incompatibility must in fact be symmetric, and if so, why?

We will offer answers to all three of these questions. For the first question—why the two reason relations of implication and incompatibility—we begin with the observation that it is not a viable explanatory strategy to begin with the bivalent semantic distinction between truth and falsity, supposed to be independently and antecedently intelligible, proceed from there to an account of the practical attitudes of doxastic acceptance and rejection (presumably parsed as taking-true and taking-false), and then somehow explain implication and incompatibility. To begin with, it is not clear how the second step is supposed to go. But in any case, we take it that one cannot understand the true/false dichotomy without understanding the propositional contents that can *be* true or false. And one cannot understand those contents without understanding the reason relations of implication and incompatibility that they stand in to one another. (This is a minimal form of semantic inferentialism.) We start instead with the pragmatics, that is, with an account of what practitioners *do*, the discursive practices they engage in. Here it seems to us that the beginning of wisdom is the realization that specifically *doxastic* acceptance and rejection are unintelligible apart from practices of contesting and defending *entitlement* to adopt those attitudes, by offering reasons for and against them. Accepting and rejecting claims and challenging and defending them with reasons come as part of an indissoluble package, no part of which is intelligible apart from the rest. Implication relations normatively govern the offering of reasons for accepting claimables and incompatibility relations normatively govern the offering of reasons for rejecting them—or (it turns out, equivalently), implications govern reasons *for* claimables and incompatibilities offer reasons *against* them. The two-fold character of reason relations, their division into implication and incompatibility relations, is an aspect of the two-fold character of practical doxastic attitudes, their division into acceptance and rejection. Truth and falsity are semantic reflections of this pragmatic reality. This pragmatics-first order of explanation will be introduced in Chapter One.

The second question is what algebraic structure material consequence relations in general should be understood to exhibit, once we have rejected the claim that they have the same structure as topological or explicitation closure operators that specifically *logical* consequence relations exhibit. We offer a detailed, systematic, constructive answer to this question. It aims to do for *material* consequence what Tarski and Gentzen did for *logical* consequence. We take this to be one of the cardinal achievements of this work. In fact, in keeping with our general strategy of understanding the reason relations of implication and incompatibility as inseparable and complementary—as two sides of one coin—we offer *one* answer to the question about the structure of which both are aspects. The conceptual and argumentative raw materials for explaining and justifying the answer we will give are assembled and developed in Chapters Two, Three, and Four. Chapter Five then defines and articulates the fundamental structure of implication and incompatibility. Although it will only at that point be intelligible, the short form of the answer is that rather than having a topological or explicitation closure structure, the open structure of reason relations has the form of a pair consisting of a commutative monoid and a distinguished implicational subspace of the space the monoid is defined on. For now this characterization is a placeholder slogan and a promissory note. But what it is a promissory note *for* is a specification of nothing less than the structure of reasons relations—and in that sense, the structure of reason—*as such*. Propositional contents should be understood as what play specific roles in reason relations of implication and incompatibility that exhibit this structure. Logic, we claim, is a set of tools for the explicit expression of such open-structured reason relations (as well as the traditionally studied, closed ones, of course). We will present (in Chapter Two) a particular logic that is provably adequate to this purpose, and in subsequent chapters investigate multiple semantic and pragmatic perspectives on the reason relations it explicates.

III. Plan of the Book

The final section of this *Introduction* offers a slightly more detailed outline of how our argument proceeds in the rest of the book.

Chapter 1: Normative Pragmatics
[Brandom]

The Introduction began with the idea that the principal task of the philosophy of logic is to explain the relationship between logic and practices of reasoning. Reason *relations* of implication and incompatibility were introduced as mediating between reasoning practices and logic. The logical expressivist order of explanation was introduced as demarcating logical concepts and the vocabulary that express them by their distinctive expressive role: making explicit, in an extension of the nonlogical language, the reason relations that govern antecedent reasoning practices in that material, prelogical language (as well as the reason relations that govern the logically extended language). The focus of the Introduction was two-fold: outlining the expressive relations between logic and prelogical material reason relations, and arguing that the structure of material relations of implication and incompatibility, which it is the job of logic to make explicit, should be understood to be structurally open, rather than having the closure structure Tarski and Gentzen had insisted on for specifically *logical* reason relations. In particular, the relations of material consequence and incompatibility logical vocabulary aims to codify are not in general monotonic, or even cautiously monotonic, and material consequence relations are not in general transitive, or cumulatively transitive.

In Chapter One, we turn from looking at the relations between reason relations of implication and incompatibility and the logical vocabulary introduced to express them in the antecedent material object language, and addresses the relations between those reason *relations* and reasoning *practices*. Thinking of semantics as the study of the conceptual contents expressed by the use (to begin with) of prelogical vocabulary and pragmatics broadly as the study of that use, we distinguish between traditional semantics-first explanations and the sort of pragmatics-first explanations of the relationship between them that we pursue in this volume. This order of explanation treats the distinction between speech acts of asserting and denying, manifesting practical attitudes of accepting and rejecting, as primitive and prior to semantic notions of truth and falsity. Both are renderings of a fundamental rational bipolarity. The link between them, which can in principle be exploited explanatorily in either direction, is that accepting is practically taking-true and rejecting is practically taking-false.

Assertion and denial evince specific kinds of acceptance and rejection. What distinguishes that doxastic species from the rest of the genus of practical attitudes pro and con is that these specifically discursive practices essentially, and not just accidentally, incorporate practices of challenging and defending the practical attitudes adopted. These are *critical* practices of asking for and giving *reasons* for those acceptances and rejections. It is argued that what it is for what is accepted or rejected to have the right sort of content—propositional, truth-evaluable, ultimately conceptual content—just is to play the right sort of role in such critical practices of demanding and providing reasons. To play that role, assertible and deniable contents must stand in two sorts of reason relations. Reasons for a claimable, that is, reasons to accept it, stand to it in relations of implication or consequences. Reasons against a claimable, that is, reasons to reject it, stand to it in relations of incompatibility. These two sorts of reason relations, implication and incompatibility, accordingly articulate the fundamental discursive bipolarity. They make visible an explanatory route from a pragmatics studying assertion and denial, to a broadly inferential semantics that understands propositional contents in terms of roles in reason relations of the two kinds.

The third part of this chapter then offers the beginnings of a normative pragmatics that might support such an inferential semantics. It begins with Greg Restall's *bilateral* normative pragmatic understanding of the turnstile of multisuccedent sequent calculi. According to it, what it means to say that Γ implies Δ (expressed in the symbolic metalanguage of the sequent calculus as "Γ|~Δ") is that the position that one would take up by asserting all the premises that are elements of the set Γ and denying all the conclusions that are elements of the set Δ would be "out of bounds": inappropriate, wrong, or forbidden. This bilateralist approach is deepened and developed by using a pragmatic metavocabulary that discerns further normative fine-structure. In the place of the single normative dimension of assessment appealed to by bilateralists, we look at two orthogonal dimensions: commitment and preclusion of entitlement. Specifying reasoning practices of challenging and defending doxastic attitudes in these more fine-grained terms permits more powerful and flexible accounts of the reason relations of implication (being a reason for) and incompatibility (being a reason against).

The three parts of Chapter One accordingly are:

a) Semantic and Pragmatic Construals of a Fundamental Bipolarity
b) Reasons and Reason Relations: Symmetries and Asymmetries
c) Bilateral Pragmatic Renderings of Reason Relations

Chapter 2: Expressive Logics
[Kaplan]

Chapter Two is also in three parts. The first part introduces the systems NM-MS and NM-SS, of NonMonotonic, MultiSuccedent logic, and NonMonotonic, Single Succedent logic, respectively. These are Gentzen-style sequent calculi. Their connective definitions conservatively extend structurally open material reason relations of implication and incompatibility. They are represented metalinguistically by triples of a prelogical language (finite set of logically atomic sentences), a set of sequents (multisuccedent or single succedent, as the case may be), and a set of materially incoherent sets of sentences. We call these structures material semantic frames (MSFs). Using the sequents of MSFs—which exhibit open, substructural reason relations—as nonlogical axioms (proof-tree leaves), instead of the reflexivity schema (as is typically done) defines reason relations, in the form of MSFs, for logically extended languages. The logic that results is thoroughly well-behaved despite dispensing with the typical structural rules.

Remarkably, the connective-rules of NM-MS are equivalent to Gentzen's rules for his version of classical logic, LK—in the context of Gentzen's strong structural closure principles. In our open, substructural context, the two sets of connective rules diverge in their behavior. NM-MS essentially uses Ketonen's fully reversible version of Gentzen's rules. NM-MS is accordingly supraclassical (and NM-SS is supraintuitionistic). If the rules of NM-MS are fed a "flat" MSF—one that consists entirely of instances of CO—the result is just classical logic, even without the stipulation of structural closure requirements. If the rules of NM-SS are fed a "flat" MSF—one that consists entirely of instances of CO—the result is just intuitionistic logic, even without the stipulation of structural closure requirements. And although the reason relations of the logically extended language are in general open and substructural (in particular, nonmonotonic and nontransitive), as they must be to constitute a conservative extension of substructural base MSFs, the purely *logical* consequence and incompatibility relations are fully structural: monotonic and transitive. These are the consequences and incompatibilities that hold, and continue to hold upon arbitrary substitution of nonlogical for nonlogical vocabulary. We present a logic for nonmonotonic consequence relations (and reason relations, more generally), that is *not* itself a nonmonotonic *logic*.

The second part of Chapter Two proves a powerful and general representation theorem for NM-MS, relating the reason relations in the logically extended language to those of the material base. It shows how to compute, for any set of sequences in the logically extended language, exactly what sequents have to hold in the underlying MSF for those sequents in the extended language to hold. Where the first part of the chapter showed how the reason relations of the logically extended language (and so the conceptual contents expressed by its vocabulary) can be conservatively *elaborated from* the reason relations governing the underlying material base vocabulary, the representation theorem makes precise the sense in which the logical vocabulary *explicitly expresses* those underlying reason relations (as well as those governing the logical extension). Together these results show exactly how and in what sense our logical systems are LX for (elaborated from and explicative of) any set of reason relations of implication and incompatibility, whether closed or open.

The third part of Chapter Two extends the expressive power of logical vocabulary beyond codifying reason relations of implication and incompatibility to making explicit regions of the MSF where *structural* conditions that do *not* hold *globally do* hold *locally*. Modal operators are introduced to mark regions of local monotonicity, transitivity, and more (e.g. classicality). The monotonicity box is most closely analogous to traditional alethic modal operators, since it marks regions of the material MSF where implications are maximally subjunctively robust: indefeasible by the addition of any collateral premises. The expressivist policy is not to make the Procrustean demand of full global structural closure on material reason relations, but to allow structural openeness and then explicitly mark local regions of the reason relations that exhibit greater structurality.

The three parts of Chapter Two are accordingly:

a) Introduction of the systems NM-MS and NM-SS, of NonMonotonic, MultiSuccedent logic, and NonMonotonic, Single Succedent logic, respectively.
b) Proof of a powerful and general representation theorem for NM-MS, relating the reason relations in the logically extended language to those of the material base.
c) Extension of the expressive power of logical vocabulary beyond codifying reason relations of implication and incompatibility to making explicit regions of the MSF where *structural* conditions that do *not* hold *globally do* hold *locally*. Modal operators are introduced to mark regions of local monotonicity, transitivity, and more.

Chapter 3. Truth-Taking and Truth-Making as Isomorphic

[Hlobil]

We started in Chapter One with a contrast between two orders of explanation, depending on whether we think of semantics or pragmatics being primary. We then developed a version of the order of explanation that takes pragmatics to come first for the special case of logical vocabulary. On the account we have developed, the characteristic expressive job of logical vocabulary is to make explicit reason relations, relations of consequence and incompatibility. That is the central claim of logical expressivism. And we have argued that while the reason relations that logical vocabulary can make explicit should not be restricted to closed consequence and incompatibility relations. Rather, the relations that logical vocabulary allows us to make explicit should include defeasible relations in which explicitation is not inconsequential. In Chapter Two, we have developed these ideas in a formal way by thinking of Material Semantic Frames, which we can represent technically as a consequence relation over a pre-logical language, i.e., a language with just atomic sentences. We saw how this allows us to recapture classical and intuitionistic logic, and how we can introduce operators that make explicit where structural conditions, like monotonicity, hold locally.

Chapter Three asks how what we have done looks from the perspective of the other order of explanation: Could we make sense of what we did so far from the perspective of a semantics-to-pragmatics order of explanation? The answer is that the formal systems that we developed in Chapter Two can be reconstructed in a pleasingly analogous way within truth-maker theory, as it has been developed by Kit Fine. In particular, the norms that we suggest govern assertions and denials – and when one is precluded from being entitled to a collection to assertions – have a mirror image in the principles that govern how worldly states make sentences true or false, and when such states are impossible.

We suggest a novel way to understand consequence in truth-maker theory. In particular, we can think of the claim "$\Gamma|\sim\Delta$" as saying that there isn't any possible state that makes every sentence in Γ and every sentence in Δ false. In the terms of our normative understanding of consequence this means that if commitment to everything in Γ precludes one from being entitled to reject everything in Δ, then this means – in metaphysical terms – that if a state makes everything in Γ true, it precludes that everything in Δ is made false. Once we understand consequence in this way in truth-maker theory, three principles that Fine often imposes on the structure of possible states emerge as the three structural principles on which we focused in Chapters One and Two, namely MO, CO, and CT.

The semantics clauses for the logical connectives in truth-makers theory can now be seen as equivalent to the Ketonen sequent rules that we used in Chapter Two. Putting together our understanding of the structural rules and the operational rules within truth-maker theory, we can recover the consequence relations of NM-MS within truth-maker theory. In particular, we can show how every consequence relation that can be codified in an NM-MS system over a material semantic frame can also be codified in truth-maker theory, and vice versa (Brandom's translation result). It also becomes easy to tweak our pragmatic-normative characterizations of consequence relations and the semantic-alethic characterizations in parallel ways. To illustrate this, we show how our correspondence between sequent calculi and truth-maker theory allows us to formulate well-known responses to the semantic paradoxes within truth-maker theory

The upshot of Chapter Three is that the pragmatics-to-semantics order of explanation that we have developed shares its structure with a version of the semantics-to-pragmatics explanatory project, namely truth-maker theory (as we developed it here, drawing on Fine's work). Hence, there is a clear sense in which both kinds of theories can codify the same reason relations. This suggests that we should study the common structure of the two frameworks in its own right, and that is what we will do in the next chapter.

Four Parts of Chapter Three:

a) Introduction of truth-maker theory as a competing, semantics-to-pragmatics account.
b) Mapping truth-maker theory and our normative pragmatic account of consequence into each other at a general and abstract level.

c) Showing how the structural rules of sequent-calculi correspond to Fine's constraints on possible states, and how operational rules correspond to Fine's semantic clauses for the connectives. Putting this together to get a truth-maker formulation of NM-MS
d) Demonstration of the flexibility of this approach by formulating the non-transitive response to the semantic paradoxes in truth-maker theory.

A short version of this material is available as Hlobil's paper "Truth-Makers and More." A longer version (including some material from Chapter Five) is available as "Truth-Maker Theory and Bilateralism."

Chapter 4: Inferential Semantics
[Kaplan]

Chapter Four has two parts. The first part introduces, develops, and deploys a distinctively inferentialist *implicational phase-space semantics* and proves soundness and completeness results for NM-MS with respect to that semantics. Inspired by Girard's phase space semantics for linear logic, this new kind of semantics offers substantial technical advances and moves far beyond it conceptually. Semantic interpretants are drawn from a space of candidate implications: pairs of premise-sets and conclusion-sets. They are "candidate" implications because only a distinguished subset of them (designated **I**) are *good* implications: implications that actually hold. (To represent incompatibilities, incoherent sets are coded as implications with empty second elements (right-hand sides), as in multisuccedent sequent calculi.) As one might expect in an inferentialist semantics, semantic interpretants are sets of candidate implications. A decisive innovation of Kaplan's implicational phase-space semantics is that it proceeds in two stages. In the first stage, the semantic interpretants are not (as one might expect) *sentences*. Rather, they are themselves (candidate) *implications* (and sets of them). Each implication is assigned a set of implications as its "ν-set." Only at the second stage of semantic interpretation will implicational interpretants be assigned to individual sentences. In a real sense, in the implicational phase-space semantics, what is in the first instance interpret*ed*, no less than the semantic interpret*ants* assigned to them, is implications rather than sentences.

What does it mean for a (candidate) implication to be assigned a set of (candidate) implications as its semantic interpretant? The implications in the ν-set of an implication are the inferentialist equivalent of its truth conditions. They are its *implicational goodness conditions*. They represent what it would take for a candidate implication to be (or become) a *good* one, one such that its conclusion actually follows from its premises. If the implication being interpreted is already a good one (in the metavocabulary of the implicational phase-space semantics, it is already in the set **I**), then its ν-set represents its *range of subjunctive robustness*: the circumstances or conditions under which it *remains* good. The guiding idea is easiest to approach by focusing exclusively on the premise side of the candidate implications that are being interpreted, and that make up the ν-sets used to interpret them semantically. Suppose that $\langle\Gamma,\Delta\rangle$ is the implication being interpreted, and that it is a good one (so $\langle\Gamma,\Delta\rangle\in\mathbf{I}$). Then for $\langle\Pi, \emptyset\rangle$ to be in the ν-set of $\langle\Gamma,\Delta\rangle$ is for $\langle\Gamma\cup\Pi,\Delta\rangle\in\mathbf{I}$. That is, $\Gamma|\sim\Delta$, and adding Π to Γ as collateral premises does not infirm or defeat the implication. Π is in the range of subjunctive robustness of the implication $\Gamma|\sim\Delta$. If $\langle\Gamma,\Delta\rangle$ is *not* already a good implication, for $\langle\Pi, \emptyset\rangle$ to be in the ν-set of $\langle\Gamma,\Delta\rangle$ is for it to be among the things that, if added to $\langle\Gamma,\Delta\rangle$, would *make* it a good implication. That is to say, as before, that $\langle\Gamma\cup\Pi,\Delta\rangle\in\mathbf{I}$. Adding the Π to Γ as auxiliary hypotheses turns the non-implication of Δ by Γ into the genuine implication $\Gamma\cup\Pi|\sim\Delta$. In general, the candidate implications in the ν-set of a candidate implication are those that, when "added to" it (their premises unioned with its premises and their conclusions unioned with its premises) yield a good implication. They are its good-makers.

In the second phase of semantic interpretation, then, inferential roles can be assigned to individual sentences. The inferential role of a sentence A is just the ordered pair of the set of good implications in which it appears as a premise and the set of good implications in which it appears as a conclusion. Those are just the ν-sets of $\langle A,\emptyset\rangle$ and $\langle\emptyset,A\rangle$, respectively. It turns out that the connective rules of the system NM-MS introduced in Chapter Two correspond to natural operations on these ν-sets. Using those operations, Kaplan is able to prove the soundness and completeness of the implicational phase-space semantics for that logical system.

So smoothly do those operations on ν-sets work with the sequent-calculus connective definitions that those soundness and completeness results are completely independent of any global structural restrictions put on the set **I** of good implications in the specification of the implicational phase-space on which logically complex sentences are semantically interpreted. In particular, **I** need not be monotonic, in that if $\langle\Gamma,\Delta\rangle\in\mathbf{I}$, then for any X $\langle\Gamma\cup X,\Delta\rangle\in\mathbf{I}$. Nor need it satisfy CT, or CM. It was shown in Chapter Two that the rules of NM-MS are basically Ketonen's reversible versions of Gentzen's sequent rules LK for classical logic. Though the equivalence breaks down in open, substructural contexts, NM-MS remains supraclassical, and if instead of applying the rules to a substantive material semantic frame one applies it only to instances of RE (of the form $A|\sim A$), the result is still wholly classical. NM-MS remains wellbehaved, not only proof-theoretically, but also semantically even in entirely open, radically substructural contexts. As we saw in Chapter Two, it codifies explicitly nonmonotonic implication and incompatibility relations and nontransitive consequence relations. Its purely logical fragment, however, is monotonic and transitive, just as the tradition requires it to be. It is a logic for nonmonotonic (and nontransitive) reason relations, rather than a nonmonotonic (or nontransitive) *logic*.

The second part uses the implicational phase-space semantics to introduce a further novel formal and conceptual apparatus for thinking about inferential roles. In the substructural context in which we are working, we show that the relationship between inferential roles—those things which the first part of the chapter provided an account of—and the implicational relationship between sentences need not correspond. That is, the following statements need not be equivalent:

1. A implies B
2. Anything that implies A also implies B:
   G |~ A, thus G|~ B
3. Anything that B implies, A also implies
   B |~ G, thus A |~ G

In fact, we can see that thinking of consequence in terms of closure forces one to run together the sort of reasoning that occurs between sentences and that which we can make about the roles that are constitutive of the meaning of those sentences. Absent closure constraints, we have room to explore the idea of *inferential role entailments,* that is the relationship that occurs between distinct inferential roles. The second point above is expressed in terms of the idea that A's role as a conclusion entails B's role as a conclusion, and *mutatis mutandis* with their roles as premises in the third point. In the formal apparatus developed this is expressed as:

$A^C => B^C$

These notions are definable in straightforward ways from the semantic framework introduced in the first half of the chapter.

The idea of an inferential role entailment is more than a technical curiosity, however. It allows us to shed light on several related phenomena. For example, we show that reflexivity connects the second and third points above with first point. That is, reflexivity of consequence allows inferential role entailments to be reflected in consequence. So, whenever $A^C => B^C$, we also have A |~ B. Transitivity allows the converse reflection (of consequence into inferential role entailments). So whenever A |~ B, we have $A^C => B^C$.

One interesting application of this framework takes inferential role entailments to provide an account of meta-inferences, a topic of some interest in the literature on substructural approaches to semantic paradoxes. Meta-inferences have attracted attention in substructural logics because, in a substructural setting, the inferences of a logic and the meta-inferences of that logic might differ (something which we are able to shed light on, as detailed in the previous paragraph). For example, ST—a non-transitive logic—looks a lot like Graham Priest's LP (logic of paradox).[32] From another vantage point, ST exhibits K3-like behavior (Kleene's strong Kleene logic). As is well known, LP and K3 are duals of each other. What is less understood is the relationship between these and logics like ST (as well as the related non-reflexive logic TS). Inferential role entailments provides a unified framework for understanding the relationship between all four of these logics (including a precise characterization of the relationship between

[32] [Reference to Priest; some think they are the same logic; references to Barrio et. Al.]

ST and TS: that they are "meta-dual").[33] Thus, "meta-inferences", as they have been characterized in the literature, show up as special cases of inferential role entailments.

Three Parts:

a) Implicational Phase-Space Semantics: which provides a tractable, technical apparatus for understanding meaning in terms of contribution to good implication.
b) Inferential Role Entailment: a logic which captures the relationship between inferential roles.
c) Application of inferential role entailment to work on meta-inferences in philosophical logic.

Chapter 5. "In the Beginning Was the Deed": The Primacy of Pragmatics Recovered

[Hlobil]

In Chapter Four, we have presented a formal account of the structure that is shared between our normative-pragmatic approach to consequence relations and the alethic-normative approach, as we find it in truth-maker theory. Chapter Five aims to explain the philosophical significance of these rules and to deepen them further.

In the first section, we argue that apparent structure in truth-maker theory that goes beyond the common structure we have identified is either illusory or should be rejected: One may wonder whether the common structure that we have identified in the previous chapter isn't leaving out what is most crucial about the semantics-first order of explanation, namely the ideas of states of the world that we represent in discourse. In particular, one might think that truth-maker theorists have offered accounts of philosophically important notions that cannot be mirrored in the pragmatics-first account, such as notions of factual equivalence, entailment, containment, or subject matter. These are the notions that articulate content according to truth-maker theorists. We show how (variants of) these notions can be understood as articulating the inferential roles of sentences in a fine-grained way. These notions turn out to be fine-grained ways to capture inferential-role entailments, as we encountered it in the cases of K3 and LP in Chapter Four.

We acknowledge that there is some structure in truth-maker theory that is not present in our phase semantics. This is exactly that structure of states that doesn't play any role in determining consequence relations. This structure is ruled out in our phase semantics by insisting on

[33] Meta-dual is a kind of meta-logical negation. It has affinities with the "conflation" operator in bi-lattices (see Fitting (citation)).

inferential roles being proper, i.e., their semantic interpretations being closed. Since this structure that we are thus ruling out doesn't contribute anything to the specification of reason relations, its inclusion in a logical theory, whose job it is to make explicit reason relations, should be avoided. Hence, the structure that we have identified has indeed a legitimate claim to be the structure of reason relations that shows up on both approaches, the semantics-first and the pragmatics-first approach.

In the second section, we explain what the structure of reason relations that is captured in our phase semantics means in philosophical terms. Mathematically speaking, this common structure is a commutative monoid with a distinguished subset. More specifically, we have a set of implications, of the form Γ|~Δ, and we can combine two such implications by combining both sides (by set-union, or fusion in the case of truth-makers), and the empty implication serves as an identity element. Since the way in which we combine the two sides of implications doesn't care about the order of the elements, our monoid is commutative. The good implications are those in our distinguished subset. But what does that mean for the structure reason relations?

The commutativity of the monoid is directly tied to the fact that neither the order of the premises nor the order of the conclusions matter. It corresponds to the traditional structural rule of permutation. Since a set, Γ, of sentences is incoherent just in case Γ|~∅, this immediately entails that incompatibility is symmetric. By contrast, consequence is not symmetric because Γ|~Δ can be in our distinguished subset of implications while Δ|~Γ is not in that subset. So it is the fact that we deal with ordered pairs as the elements in our monoid set that ensures that consequence is not symmetric, while it is the commutativity of our monoidal operational that ensures that incompatibility is symmetric. What we see here is that the bilateral structure of reasons --- assertion and denial ---, which corresponds to the two elements in our ordered pairs, is the sources of the asymmetry of consequence. And the order-independence of combinations of assertions and denials is what corresponds to the symmetry of incompatibility. Thus, reasons are structured by being able to play two kinds of roles, where their combination within either role is order-independent.

What it means for reason relations to be open is that they don't have any significant structure that doesn't arise out of this structure of two kinds of roles, each of which is order-independent, including the way in which particular vocabulary can interact with this structure. More specifically, note that the distinguished subset of our monoid set swings free from the monoidal operation (in contrast to what we would see in a residuated lattice). It is this independence that allows us to reject monotonicity and transitivity as global structural constraints on reason relations. In truth-maker theory, this independence shows up as the independence of the algebra of states from their possibility. Local regions in which a certain structure holds are either a brute fact or they are ensured by some particular vocabulary in that region, such as logical vocabulary.

In the third section, we argue that the structure of reason relations that we found is what we should expect from the perspective of logical expressivism. We start with the traditional structural rules, weakening and cut. The independence of the monoidal operation and the distinguished subset allows us to codify open reason relations; i.e., what allows us to codify such open relations is that the mere fact that an implication is the result of applying our monoidal operation to two other implications with a particular status doesn't necessarily tell us anything about the status of our target implication. That is what allows us to say that the badness of an implication can be healed by adding assertions or denials, and it allows us to say that while two implications are both good, we cannot always string them together to get another good implication. It allows us to codify reason relations in which explicitation is not always inconsequential. Since we have good philosophical reasons to think that expliciation can be consequential, we have good reasons to accept this independence of our monoidal operation and our distinguished subset of the monoid set.

What about the symmetry of incompatibility? In our phase semantics, this shows up as the commutativity of our monoid, or equivalently as the order-independence of premises and conclusions. So one way to argue for this is that our commitments and rejections are on a par, in the sense that none of them comes first or second, etc. While this is indeed plausible, our ROLE colleague Ryan Simonelli has presented an additional consideration: It would undermine the normative role that reasons against claims are playing to allow that incompatibility is asymmetric. For it is essential for something being an act of giving a reason against a claim that the person whose commitment is thus questioned cannot simply drop the commitment in question, take the reason against on board, and then add the controversial commitment again. We will return to the additional structure that consists in the division between assertions and denials in the next chapter.

What makes this structure particularly plausible from the perspective of logical expressivism is that it allows us to codify open reason relations while also giving us the structure needed for reasons to play their essential roles, such as the symmetry of incompatibility that is required for something to play the role of a reason against. And this structure allows us to introduce vocabulary, namely logical vocabulary, that can (a) make explicit such reason relations while (b) ensuring that the traditional structure is available in the region that is created by that vocabulary.

Three Parts:

a) Recapturing truth-maker notions of content within our normative-pragmatic approach.
b) Explain what the structure of the commutative monoid plus subset means in philosophical terms.
c) Explain why the commutative monoid with a distinguished subset is the right structure, given logical expressivism.

Chapter 6: Conclusion.
[Hlobil]

We started this book by suggesting a normative-pragmatic understanding of logic as making explicit reason relations, viz. the relations that mediate between logic and good reasoning. We have rejected the claim that these reason relations must be topological closure relations, and even the idea that they are closure relations in any weaker sense. Thus, we have endorsed the possibility of open consequence relations and open incompatibility relations. We have illustrated what this means for logic by presenting various formal systems in Chapter Two. In Chapters Three and Four, we have seen that the structure that results from this approach is very general and can also be found in truth-maker theory. Although reason relations do not obey the traditional structural rules, their structure forms a commutative monoid with a distinguished subset. And we have explained, in Chapter Five, why this is the structure that we should expect reason relations to have, given logical expressivism. The symmetry of incompatibility and the asymmetry of consequence fall out of this structure, and they are what we would expect when we look at reason relations from a pragmatic-normative perspective. However, there is an aspect of this structure that we didn't discuss in the previous chapter, namely the fact that contentful elements can occur in two kinds of role: assertions and denials. This is where we started in Chapter One, with the discussion of normative bilateralism. And this is where we will end in this conclusion. But we will get to this by addressing an even larger question: Why are there two reason relations, implication and incompatibility? Why not just one, or three? And if there must be two, why just *these* two?

To answer this question we return to the normative pragmatics with which we started, and we will look in particular at the relations between reason relations and reasoning practices. What it is to have content is, in the first instance, to be a move in our discursive practice of giving and asking for reasons. This practical of giving and asking for reasons requires that we can not only give reasons for a claim but that we can also challenge a claim by giving reasons against it. The reason relations that allow something to be supported and challenged by reasons are what articulates the content to which the attitudes are attitudes. Practices in which we could give only reasons for or only reasons against claims are not practices of making claims at all. That is why we need two and not just one reason relations, implication and incompatibility. We don't need a third relation because a practice in which we can give reasons for and reasons against claims is sufficient to have a practice of claims, that is to have a practice in which actions can have contents.

We respond appropriately to good reasons for a claim by accepting the claim; and we respond appropriately to good reasons against a claim by rejecting the claim. These two attitudes are hence what is minimally needed to have a practice of giving and asking for reasons. And what it is to accept a claim is simply to do that which is the appropriate response to good reasons for the claim. And what it is to reject a claim is simply to do that which is the appropriate response to good reasons against the claim. Thus, we should understand the attitudes of acceptance and rejection – and with them the speech acts of assertion and denial – in terms of our practice of giving and asking for reasons and the reason relations that constrain that practice.

The semantic notions of truth and falsity can then be understood as the representationalist reflection of acceptance and rejection and, hence, ultimately in terms of reasons for and reasons against claims. To be true or false is to be made true or false by how things are. We understand what it is that a sentence is being made true or false by how things are in analogy to what it is for the sentence to be accepted or rejected. To think of the states that could make our assertions and denials true or false is to think of states that alethically exclude each other in a way that is isomorphic to how our assertions and denials normatively exclude each other. That was one of the upshots of the isomorphism between normative bilateralism and truth-maker theory.

To understand the contents that we are accepting and rejecting, we must understand their relations of consequence and incompatibility. For contents are articulated by the reason relations in which they occur. And in order to understand these relations, we must understand our practice of giving and asking for reasons. Hence, although the structure of contents that is captured by our commutative monoids with distinguished subsets is instantiated in the contents that we assert and deny as well as in the states that make our assertions and denials true and false, the normative-pragmatic side of this coin is primary in the order of explaining content.

We end this concluding chapter with a final discussion of logical expressivism. …

Four Parts:

a) Summarizing what we did and raising the question about why there are exactly two kinds of reason relation.
b) Giving a normative-pragmatic account of reason relations that explains why there are two such relations, and why they are implication and incompatibility.
c) Explaining the primary of pragmatics over semantics, despite the isomorphism.
d) Return to logical expressivism: What have we learned about logic? In what sense is the book a defense of logical expressivism? Etc.