

Articulating Reasons: Chapter Six

Objectivity and the Normative Fine Structure of Rationality

I

A basic pragmatist methodological thesis is that the point of the theoretical association of *meanings* with linguistic expressions is to explain the *use* of those expressions. (Semantics must answer to pragmatics.) A fundamental divide among theorists who agree in endorsing this methodological pragmatism then concerns the terms in which the use of linguistic expressions is understood. One camp takes as its explanatory target *proprieties* of use. Meanings are invoked to explain how it is *correct* or *appropriate* to use words and sentences, how one *ought* to deploy them. The other camp (Quinean behaviorists may serve as an example) insists on specifying the use to be explained in sparser terms. The ultimate explanatory target at which semantic theory aims is utterances and dispositions to utter described in a vocabulary resolutely restricted to nonnormative terms.¹ I'll say something further along about why I think the second camp is misguided.

¹ It might be noticed in passing that it is *not* harmless to paraphrase this choice as that between talking about how linguistic expressions *ought* to be used, and how they are *actually* or *in fact* used, or how practitioners are *disposed* to use them. Using an expression correctly or incorrectly is something practitioners can actually or in fact do, something they can be disposed to do. The difference should be

But for now I just want to put this option to one side, with the observation that doing so does not by itself require relinquishing commitments to naturalistic semantics. For one might well accept a normative characterization of the explanatory target—specifying use in terms that permit one to distinguish, say, correct from incorrect representations of states of affairs—while retaining a commitment eventually to offering a reductive account of the origin and nature of those proprieties in turn, framed in the sort of modally rich but not explicitly normative vocabularies routinely employed in the special sciences, whether physical, biological, or social.²

The idea behind assertibility theories of the propositional contents expressed by declarative sentences is to start with a notion of linguistic propriety that could be understood in terms of allowable moves in a game. To specify the circumstances in which a sentence is assertible is to say when its assertional use is appropriate or allowable, when a speaker is licensed or entitled to use the sentence to perform that speech act, when its assertional utterance would have a certain sort of normative significance or status. Basing one's semantics on the association of sentences with assertibility conditions is not only a way of construing meaning as potentially explanatory of use. It is an identification of meaning with a core feature of use—one, presumably, in terms of which other important dimensions of use can then be explained. The very tight connection that is envisaged between meaning, so construed, and proprieties of use is, I

located rather in the vocabulary the theorist is permitted to use in characterizing what speakers and audiences actually do and are disposed to do. Formulating this difference as a difference between saying how the language *is* used and how it (only) *ought* to be used is the decisive move in the conjuring trick that lands one in the intractable puzzles about conceptual normativity that Kripke's Wittgenstein has made familiar.

think, one of the sources of the attractiveness of broadly assertibilist approaches to meaning.

Another is the prospect of starting with relatively clear explanatory raw materials. The first obligation of the assertibility theorist will of course be to explain the notion of assertibility. Doing that requires first saying something about assertional force: about what it is for a speech act to have the of significance of an assertion. The next requirement is to specify a sense of propriety appropriate to that speech act: to say what it is for an assertion to be appropriate or correct, for the speaker to be entitled or permitted to produce it. Neither of these tasks is simple or straightforward. But we do have a relatively familiar and unmysterious framework in which to address them. For the first takes its place as an instance of distinguishing different kinds of moves in a game; we are invited to think of asserting as a species in the same genus with punting, bidding, castling, betting, and so on. And the second takes its place as an instance of saying when moves of the specified kind are permitted. We should count ourselves fortunate indeed if we could, as the assertibilist hopes and promises, construct a workable concept of the meaning or content associated with declarative sentences (and hence also with the beliefs and judgments they express) from such raw materials.

The biggest challenge to this happy prospect stems from the fact that assertions are subject to two essential, but fundamentally different kinds of normative appraisal. We can ask whether an assertion is correct in the sense that the speaker was entitled to make

² I think of Dretske, Fodor, and Millikan as presenting theories with this general shape. Perhaps Gibbard's very different approach to moral norms, when generalized and adapted to the case of linguistic

it, perhaps in virtue of having reasons, evidence, or some other sort of justification for it. This might be thought of as a way of asking whether the speaker is blameworthy for performing this speech act, whether the speaker has fulfilled the obligations the rules of the game specify as preconditions for making a move of this sort in the game. This is the normative aspect of use the assertibilist begins with. But we can also ask whether the assertion is correct in the sense of being *true*, in the sense that things are as it claims that they are. It is a basic criterion of adequacy of a semantic theory that it explain this dimension of normative assessment, this normatively described aspect of use. The challenge to the sort of approach to semantics I have been calling ‘assertibilist’ is to show how the conceptual raw materials this approach allows itself can be deployed so as to underwrite attributions of propositional content for which this sort of *objective* normative assessment is intelligible.

The attempt by assertibility theorists to satisfy this central criterion of adequacy of semantic theories has typically taken the form of appeals to some sort of *ideality* condition. Assessments of truth are understood as assessments of assertibility under ideal conditions (what Sellars called ‘semantic assertibility’)—of what claims one would be entitled to or justified in making if one were an ideal knower, or given full information, maximal evidence, at the end of inquiry, and so on. I’m not going to argue the point here, but my own view is that this sort of strategy is hopeless.³ If it is the best available, we

norms, will find its place here too.

³ My thought is that there is no way to specify the ideality in question that is not either question-begging (in implicitly appealing to a notion of truth) or trivial, in the light of the sensitivity of the practical effects of more ideal status for one belief both to the falsity of collateral beliefs, and even to ignorance concerning them. I present one argument along these lines in "Unsuccessful Semantics" *Analysis* Vol. 54 No. 3 (July 1994) pp. 175-8.

should just give up the assertibilist project. In that case the obvious alternative is to start with a notion of meaning that directly underwrites normative assessments of objective representational correctness: truth conditions. We will not then be able to explain the association with linguistic expressions of semantic contents, so understood, by straightforward assimilation to making moves allowed by the rules defining a game, as promised by the alternative broadly assertibilist explanatory strategy. Attempts by truth conditional semantic theorists to construct the other dimension of normative assessment of assertions—assertibility in the sense of entitlement, justification, having reasons or evidence—have typically taken the form of reliability theories. Assessments of assertibility in the sense of cognitive entitlement or justification are understood as assessments of objective or subjective likelihood of truth. In Chapter Three I rehearsed some of the structural problems afflicting this sort of strategy as well.

What I want to do instead is to explore a different way in which one might start from the sort of normative statuses the assertibilist invokes, intelligible in terms of moves in a rule-governed game, and on that basis associate with declarative sentences propositional contents that are objective in the sense of swinging free of the attitudes of the linguistic practitioners who deploy them in assertions. The idea is roughly to split up the notion of assertibility into two parts. More precisely, where assertibility theorists appeal to just *one* sort of normative status—a sentence being assertible, or a speaker being justified or having sufficient reasons to assert it—I'll look at *two* kinds of normative status: commitment and entitlement. Discerning this additional normative structure in linguistic practice, in particular, exploiting the relations and interactions between these two kinds of

normative status articulating the *force* or significance of linguistic performances, makes possible the specification of propositional *contents* with desirable properties. Chief among these is the *objectivity*, in the sense of a specifiable sort of attitude-transcendence, of the propositional contents that are suitably defined in terms of the roles played by their bearers in linguistic practices characterized in terms of alterations and inheritance of commitments and entitlements. This result holds good even if the normative statuses of commitment and entitlement are themselves understood as *social* statuses, that is, as creatures of individual and communal attitudes.

II

Semantic assertibilism is implicitly committed to demarcating specifically *linguistic* practices by restricting that term to practices that confer on some performances the significance of *claims* or *assertions*. What is asserted in an act of asserting, what is assertible, is a propositional content. Assertible contents, assertibles, are also believables and judgeables; states of belief and acts of judgment can accordingly be expressed by assertions. Linguistic expressions whose free-standing utterances have the default significance of assertions are (declarative) sentences. Our aim is to investigate the propositional contents that are associated with linguistic expressions by their playing this central role in assertional practices.

The first key idea is that a performance deserves to count as having the significance of an assertion only in the context of a set of social practices with the structure of (in Sellars's phrase) a *game of giving and asking for reasons*. Assertions are essentially performances that can both serve as and stand in need of reasons. Propositional contents are essentially what can serve as both premises and conclusions of inferences. This inferentialist idea might be called "linguistic rationalism."⁴ Linguistic rationalism is not a standard part of the armamentarium of semantic assertibilism, but I think it is what is required to make that explanatory strategy work. I suggested in Chapter One what seem to me good reasons to see giving and asking for reasons as the defining core of discursive (concept-mongering) practice; I do not propose to rehearse them here. Rather, I want to treat linguistic rationalism as a hypothesis, and to explore its consequences.

In the rest of this chapter, I want to make two arguments. First, in this section I will argue that no set of practices is recognizable as a game of giving and asking for reasons for assertions unless it involves acknowledging at least two sorts of normative status, *commitments* and *entitlements*, and some general structures relating them. I'll show how

⁴ It is not identical with inferentialism as introduced in Chapter 1, since that thesis concerned the relative explanatory priority of the concepts of inference and representation, and linguistic rationalism as used here is silent about representation. In the Introduction I distinguished three sorts of inferentialist claims: weak inferentialism, strong inferentialism, and hyperinferentialism. Weak inferentialism is the claim that inferential articulation is a *necessary* aspect of conceptual content. Strong inferentialism is the claim that *broadly* inferential articulation is *sufficient* to determine conceptual content (including its representational dimension). Hyperinferentialism is the claim that *narrowly* inferential articulation is *sufficient* to determine conceptual content. Broadly inferential articulation includes as inferential the relation even between circumstances and consequences of application, even when one or the other is *noninferential* (as with observable and immediately practical concepts), since in applying any concept one implicitly endorses the propriety of the inference from its circumstances to its consequences of application. Narrowly inferential articulation is restricted to what Sellars calls "language-language" moves, that is, to the relation between propositional contents. Weak inferentialism is the most plausible of these theses. Strong inferentialism is the view endorsed and defended here and in *MIE*. Hyperinferentialism is plausible at most for some abstract mathematical concepts. Linguistic rationalism is a version of weak inferentialism, which the present chapter endeavors to show has some strong inferentialist consequences, when suitably elaborated.

we can understand practices incorporating those statuses in that structure as conferring propositional contents on linguistic expressions suitably caught up in them. Then, in the next section, I'll argue that propositional contents specified in terms of their contribution to the commitments and entitlements that articulate the normative significance of speech acts exhibiting those contents exhibit *objectivity* of a particular sort: they are not about any constellation of attitudes on the part of the linguistic practitioners who produce and consume them as reasons.

Suppose we have a set of counters or markers such that producing or playing one has the social significance of making an assertional move in the game. We can call such counters ‘sentences’. Then for any player at any time there must be a way of partitioning sentences into two classes, by distinguishing somehow those that he is disposed or otherwise prepared to assert (perhaps when suitable prompted). These counters, which are distinguished by bearing the player’s mark, being on his list, or being kept in his box, constitute his score. By playing a new counter, making an assertion, one alters one’s own score, and perhaps that of others.

Here is my first claim: for such a game or set of toy practices to be recognizable as involving assertions, it must be the case that playing one counter, or otherwise adding it to one’s score, can *commit* one to playing others, or adding them to one’s score. If one asserts “The swatch is red,” one *ought* to add to one’s score also “The swatch is colored.” Making the one move *obliges* one to be prepared to make the other as well. This is not to say that all players actually *do* have the dispositions they *ought* to have. One may not act

as one is committed or obliged to act; one can break or fail to follow this sort of rule of the game, at least in particular cases, without thereby being expelled from the company of players of the asserting game. Still, I claim, assertional games must have rules of this sort: rules of *consequential commitment*.

Why? Because to be recognizable as assertional, a move must not be idle, it must make a difference, it must have consequences for what else it is appropriate to do, according to the rules of the game. Assertions express judgments or beliefs. Putting a sentence on one's list of judgments, putting it in one's belief box, has consequences for how one ought, rationally, to act, judge, and believe. We may be able to construct cases where it is intelligible to attribute beliefs that are consequentially inert and isolated from their fellows: "I just believe that cows look goofy, that's all. Nothing follows from that, and I am not obliged to act in any particular way on that belief." But *all* of our beliefs could not intelligibly be understood to be like this. If putting sentences onto my list or into my box *never* has consequences for what else belongs there, then we ought not to understand the list as consisting of all my judgments, or the box as containing all my beliefs. For in that case knowing what moves someone was disposed to make would tell us nothing else about that person.

Understanding a claim, the significance of an assertional move, requires understanding at least some of its consequences, knowing what else (what other moves) one would be committing oneself to by making that claim. A parrot, we can imagine, can produce an utterance perceptually indistinguishable from an assertion of "The swatch is red." Our

nonetheless not taking it to have asserted that sentence, not to have made a move in that game, *is* our taking it that, unaware as it is of the inferential involvements of the claim that it would be expressing, of what it would be committing itself to were it to make the claim, it has not thereby succeeded in committed itself to anything. Making that assertion is committing oneself to such consequences as that the swatch is colored, that is not green, and so on.

For this reason we can understand making a claim as taking up a particular sort of normative stance towards an inferentially articulated content. It is *endorsing* it, taking *responsibility* for it, *committing* oneself to it. The difference between treating something as a claiming and treating it just as a brute sounding off, between treating it as making a move in the assertional game and treating it as an idle performance, is just whether one treats it as the undertaking of a commitment that is suitably articulated by its consequential relations to other commitments. These are *rational* relations, whereby undertaking one commitment *rationally* obliges one to undertake others, related to it as its inferential consequences. These relations articulate the *content* of the commitment or responsibility one undertakes by asserting a sentence. Apart from such relations, there is no such content, hence no assertion.

I have been belaboring what is perhaps an obvious point. Not just any way of distinguishing some sentences from others can be understood as distinguishing those asserted, those that express judgments or beliefs from the rest. For putting a sentence on a list or in a box to be intelligible as asserting or believing it, doing so must at least have

the significance of committing or obliging one to make other moves of a similar sort, with sentences that (thereby) count as inferentially related to the original. Absent such consequential commitments, the game lacks the rational structure required for us to understand its moves as the making of contentful assertions.

The next claim I want to make is that practices incorporating a game of giving and asking for reasons—*rational* practices, which linguistic rationalism supposes to be the only ones that deserve to be thought of as *linguistic* practices—must involve acknowledgment of a *second* kind of normative status. We have said that making a move in the assertional game should be understood as acknowledging a certain sort of *commitment*, articulated by consequential inferential relations linking the asserted sentence to other sentences.

But players of the game of giving and asking for reasons must also distinguish among the commitments an interlocutor undertakes, a distinguished subclass to which she is *entitled*. Linguistic rationalism understands assertions, the fundamental sort of speech act, as essentially things that can both serve as and stand in need of reasons. Giving reasons for a claim is producing other assertions that *license* or *entitle* one to it, that *justify* it. Asking for reasons for a claim is asking for its warrant, for what entitles one to that commitment. Such a practice presupposes a distinction between assertional commitments to which one is entitled and those to which one is not entitled. Reason-giving practices make sense only if there can be an issue as to whether or not practitioners are entitled to their commitments.

Indeed, I take it that liability to demands for justification, that is demonstration of entitlement, is another major dimension of the responsibility one undertakes, the commitment one makes, in asserting something. In making an assertion one implicitly acknowledges the propriety, at least under some circumstances, of demands for reasons, for justification of the claim one has endorsed, the commitment one has undertaken. Besides the *committive* dimension of assertional practice, there is the *critical* dimension: the aspect of the practice in which the propriety of those commitments is assessed. Apart from this critical dimension, the notion of *reasons* gets no grip.

So the overall claim is that the sense of endorsement that determines the force of assertional speech acts involves, at a minimum, a kind of *commitment* the speaker's *entitlement* to which is always potentially at issue. The assertible contents expressed by declarative sentences whose utterance can have this sort of force must accordingly be inferentially articulated along both normative dimensions. Downstream, they must have inferential *consequences*, commitment to which is entailed by commitment to the original content. Upstream, they must have inferential *antecedents*, relations to contents that can serve as premises from which entitlement to the original content can be inherited.

These two flavors of normative status are not simply independent of one another. They interact. For the entitlements at issue are entitlements to commitments. We can say that two assertible contents are *incompatible* in case *commitment* to one precludes *entitlement* to the other. Thus commitment to the content expressed by the sentence "The swatch is red," rules out entitlement to the commitment that would be undertaken by asserting the

sentence “The swatch is green.” Incompatibilities among the *contents* expressed by sentences, derived from the interaction of the two normative dimensions articulating the *force* of assertions of those sentences, induce their own sort of inferential relation. For we can associate with each sentence the set of all the sentences that are incompatible with it, according to the rules of the particular assertional game of giving and asking for reasons within which it plays a role. Inclusion relations among these sets then correspond to inferential relations among the sentences. That is, the content of the claim expressed by asserting “The swatch is vermillion,” entails the content of the claim expressed by asserting “The swatch is red,” because everything incompatible with being red is incompatible with being vermillion.⁵

So the two sorts of normative status that must be in play in practices that incorporate a game of giving and asking for reasons, commitment and entitlement, induce *three* sorts of inferential relations in the assertible contents expressed by sentences suitably caught up in those practices:

- *committive* (that is, commitment preserving) inferences, a category that generalizes deductive inference,
- *permissive* (that is, entitlement preserving) inferences, a category that generalizes inductive inference, and
- *incompatibility* entailments, a category that generalizes modal (counterfactual supporting) inference.

⁵ It should be remarked that acknowledging incompatibilities means treating the assessment of entitlements as a two-stage process. First one assesses *prima facie* claims to entitlement, and then winnows from this set those commitments that are incompatible with other commitments, and hence precluded from

It can be argued on relatively general grounds, though I will not do so here, that these three sorts of inferential consequence relation can be ranked strictly by their strength: all incompatibility entailments are commitment preserving (though not vice versa) and all commitment preserving inferences are entitlement preserving (though not vice versa).

This is what in title of the chapter I call “the normative fine structure of rationality.” Rational practices, practices that include the production and consumption of reasons—the “giving and asking for reasons,” of the Sellarsian slogan with which we began—must distinguish two sorts of normative status: a kind of *commitment*, undertaken by the assertional speech acts by which alone anything can be put forward *as a reason*, and a kind of *entitlement*, which is what is at issue when a reason is requested or required. This normative fine structure is *inferentially articulated* along three axes, defined by inheritance of commitment, inheritance of entitlement, and entailments according to the incompatibilities defined by the interactions of commitments and entitlements.

The core idea behind assertibility theories was a pragmatist one. It is to start with something we *do*—specifically, to start with the fundamental speech act of *assertion*, with the notion of assertional *force*—and to read off a notion of *content* (what we say or think) directly from proprieties governing that sort of speech act. Thus the content expressed by declarative sentences was to be identified and articulated in terms of assertibility conditions: that is, conditions under which it would be *appropriate* to assert the sentence. I have suggested that in the context of a commitment to linguistic

entitlement. What I call (here and below) "entitlement preserving inferences" structure the inheritance of *prima facie* commitments.

rationalism, to the idea that the game of giving and asking for reasons is the home language game of assertion, this undifferentiated normative notion of the propriety of an assertion can be replaced by a more finely articulated normative structure. For the game of giving and asking for reasons reveals itself as involving two different sorts of normative status (and so normative assessment). The score we must keep on those who engage in practices that include giving and asking for reasons has two components; we must keep track of what they are *committed* to, and also of which of these commitments they are *entitled* to.

Making this refinement at the level of the *pragmatic* theory, the theory of assertional *force*, induces corresponding refinements at the level of *semantic* theory, the theory of assertible *content*. For now instead of the undifferentiated question “Under what circumstances would it be appropriate to assert the sentence?” we must ask “Under what circumstances (for instance, in the context of what other claims) would one count as *committed* to the claim expressed by the sentence?” and “Under what circumstances (for instance, in the context of what other claims) would one count as *entitled* to the claim?” Indeed, it appears that we should not only look upstream, by asking what claims or circumstances commit or entitle us to the claim in question, but also downstream by asking to what else the claim in question commits or entitles us as *consequences*. Further, we should take account of the interaction of these two normative dimensions into which we have subdivided the undifferentiated notion of assertibility or appropriate assertion, by asking also with what other claims the claim in question is *incompatible*. This structure gives broadly assertibilist semantic theories—those that seek to derive a

notion of semantic content directly from the proprieties of use that are the subject matter in the first instance of pragmatics—a great deal more to work with.

III

What I want to do in this final section is to demonstrate one of the semantic payoffs that this richer pragmatic structure enables.

Assertibilist semantic theories seek to understand propositional content by associating with sentences as their semantic interpretants *assertibility conditions*: circumstances under which the sentence in question is appropriately assertible. The attraction of such theories is due to the very close tie they establish between meaning and use. They hold out the promise of reading *semantic* norms directly off of *pragmatic* ones, that is, off of the rules for the asserting game, or the norms implicitly acknowledged by those who participate in assertional practice. The challenge for them is to get out the other end of their machinery a sense of ‘correct’ that is sufficiently objective to be recognizable as a notion of propositional content. On the face of it, assertional speech acts are subject to two central sorts of normative appraisal. One asks whether the speech act was appropriate in light of the attitudes of the practitioners: Was all available evidence taken into account? Were the inferences made good ones, as far as the practitioners know? In general, did the speaker follow the rules of the game, so as not to be blameworthy for producing the assertion? The other sort of appraisal swings free of the attitudes of the practitioners, and looks instead to the subject matter about which claims are made for the applicable norms. Here the central question is: Is the claim correct in the sense that

things really are as it says they are? Only an omniscient being could follow a rule that enjoining practitioners to make only claims that are true. This means that the conduct of those who, through no fault of their own, make false claims is not blameworthy. Nonetheless, this further sort of appraisal is possible.

So theories of this sort face a structural dilemma. In order to make their raw materials as intelligible as possible, one wants to tie assertibility closely to people's attitudes, to what they *take* to be assertible or *treat as* assertible. This need not take the extreme form of identifying the assertibility conditions of sentences with nonnormatively specified conditions under which practitioners are disposed to assert those sentences. But there is pressure to make whatever norms are invoked be ones that can be read off of the attitudes of practitioners who apply and acknowledge the applicability of those norms. On the other hand, the more closely the norms of assertibility that articulate the contents associated with sentences reflect the attitudes of those who use the sentences, the farther they will be from the sort of objective norms appealed to in assessments of representational correctness, of getting things right according to a standard set by the things about which one is speaking. If 'assertible' is read as requiring correctness in this more objective sense, then assertibility conditions just become truth conditions, and the link to the attitudes and practices of those who use the sentences to make claims, which promised to make the association of sentences with semantic content intelligible, becomes correspondingly obscured. So the challenge for assertibility theories is to start with a notion of propriety of assertion that is grounded in and intelligible in terms of the

practice of speakers and audiences, and yet which is rich enough to fund normative assessments that are objective in the sense of transcending the attitudes of practitioners.

Consider an example of the sort that standardly causes trouble for assertibility theories.

Whenever

- 1) “The swatch is red,”

is appropriately assertible, it is equally appropriate to assert

- 2) “The claim that the swatch is red is properly assertible by me now.”

For the latter just makes explicit, as part of the content that is asserted, what it is implicit in the what one is doing in the former asserting. And yet, we want to say that the contents are different. Though the two claims have the same *assertibility* conditions, they have different *truth* conditions. For the swatch could *be* red without me being in a position to *say* that it is. And surely we could describe circumstances in which I would have extremely good evidence that the swatch was red, so that (1) is assertible for me, even though it in fact was not red—perhaps even in circumstances where the swatch does not exist. It seems that assertibility theories are leaving out something important.

But things look different if we help ourselves to the finer-grained normative vocabulary of commitment and entitlement, and hence of incompatibility. (1) and (2) would be incompatibility equivalent (in the sense that they incompatibility-entail one another) just in case everything incompatible with (1) were incompatible with (2), and vice versa. But in the situations just described, this is precisely not so. To say that the swatch could be red without me being in a position to say that it is is to say that some claims are

incompatible with (1) being assertible by me now that are not incompatible with (1). For instance,

- 3) “I do not exist,” or
- 4) “Rational beings never evolved,”

are both incompatible with (2), but not with (1).

And to say that there are circumstances in which I would have extremely good evidence that (1) is true, so that it is appropriately assertible by me, even though (1) is not in fact true is just to say there are claims that are incompatible with (1), but not with its being assertible by me.

- 5) “In the absence of a swatch, but otherwise in circumstances that are perceptually quite standard, my optic nerve is being stimulated just as it would be if there were a red swatch in front of me,”

might qualify. The additional normative expressive resources made available by distinguishing the status of being assertionally *committed* from that of being *entitled* to such a commitment are sufficient to distinguish the contents of ordinary claims from those of claims about what is assertible.

One might worry that this result is not robust, but depends on setting up the test case in terms of the undifferentiated notion of appropriate assertibility, while assessing it using the more specific normative notions of commitment and entitlement (and so incompatibility). This thought suggests that better test cases would be provided by:

- 2') “I am now committed to the claim that the swatch is red,” and
- 2'') “I am now entitled to the claim that the swatch is red.”

But in fact this additional specificity makes no difference. (3) and (4) are incompatible with both (2') and (2''), just as they were with (2), though not incompatible with (1). And (5), or some variant of it, is still incompatible with (1), but not with (2') or (2'').

In fact, looking at (2') and (2'') offers some insight into *why* distinguishing the normative statuses of commitment and entitlement offers an important expressive advance in broadly assertibilist semantic theories, when compared with the vaguer notion of assertibility. For although one is *committed* to (2') whenever one is *committed* to (1), one is not *entitled* to those claims in all the same circumstances. In particular, I can be *entitled* to (2') just on the basis of a rehearsal of my commitments, perhaps by noticing that I just asserted (1), without needing to investigate the colors of swatches. But I can only become *entitled* to (1) by an investigation of just that sort. In the other case, it is not at all clear even that one is *entitled* to (2'') whenever one is entitled to (1). Insofar as reliabilism is correct (what I called the “Founding Insight” of reliabilism, in Chapter Three), I can *be* entitled to claims without having good reason to *believe* that I am so entitled. But even if that is wrong, and entitlements to claims of the form of (2'') do go along with entitlements to base-level claims such as (1), the two sorts of claims are still distinguishable in terms of the *commitments* they involve. For surely one could be *committed* to the claim that the swatch is red, that is to (1), without thereby being committed to the claim that one is *entitled* to it. In general one *ought* to be entitled to one’s commitments, but the game of giving and asking for reasons has a point precisely insofar as we must distinguish between commitments to which one is entitled and those to which one is not. So one must at least allow that it is *possible* that one is in such a

situation in any particular case. Again, (2'') and (1) do not evidently have the same commitment-inferential *consequences*. The conditional:

6) “If the swatch is red, then the swatch is red,”

is evidently correct in that it codifies a commitment-preserving inference. (The stuttering inference is as safe as any could be.) By contrast, the conditional

7) “If I am entitled to the claim that the swatch is red, then the swatch is red,”

is not one that ought to be endorsed as correct in the sense of commitment-preserving, at least for any notion of entitlement that humans can secure regarding empirical matters of fact. It is, after all, an instance of the very implausible schema:

8) “If S is entitled to the claim that the swatch is red, then the swatch is red.”

Now I have been careful to be as noncommittal as possible regarding the specifics of the notions of commitment and entitlement (and hence incompatibility) employed in discussing these examples. For that reason, some of my particular claims about what are and are not good inferences, in any of the three fundamental senses of the permissive, committive, or incompatibility entailments, will be controversial for those who have in mind some particular ways of thinking about commitment and (especially) entitlement. But worries about these details will not affect the overall point I am after. For that is that notions of commitment and entitlement (and hence of incompatibility) *can* be put in play so as rigorously and systematically to distinguish between the contents of ordinary empirical claims and the contents of any claims about who is committed or entitled to what. The fact that other ways of deploying the notions of commitment and entitlement

would *not* allow all of those distinctions is neither here nor there; it would just provide a good reason not to use *those* notions of commitment and entitlement.

The fact is that the distinction between sentences sharing *assertibility* conditions and sharing *truth* conditions, illustrated for instance by sentences such as:

- 9) “I will write a book about Hegel,” and
- 10) “I foresee that I will write a book about Hegel,”

which are alike in the first way, but not in the second, can be made out in terms of commitments and entitlements, without the need to invoke the notion of truth. I may be *committed* to (9) and (10) in the same circumstances, and may even be *entitled* to them in the same circumstances; we could regiment the use of ‘foresee’ so as to ensure this. But

- 11) “I will die in the next ten minutes,”

will still be incompatible with (9) and not with (10), for any notion of foreseeing that does not entail omniscience.⁶ And we should not be surprised by this result. For the *consequences* of (9) and (10) are quite different.

- 12) “If I will write a book about Hegel, then I will write a book about Hegel,” is, once again, as secure an inference as one could wish.

- 13) “If I foresee that I will write a book about Hegel, then I will write a book about Hegel,”

⁶ As Crispin Wright has pointed out, according to the definitions offered here, if two claims differ in their incompatibilities, they can at most be alike in the circumstances in which one is *prima facie* entitled to them, not in the circumstances in which one is *finally* entitled to them. The assertibilist tradition did not make this distinction, since it did not divide the undifferentiated status of assertibility into commitment and entitlement in the first place (and hence was not in a position to discuss incompatibility). I think a good case can be made for treating the bits of their motivations that (implicitly) concern entitlement as addressing *prima facie* entitlements.

by contrast, is a conditional whose plausibility depends on how good I am at foreseeing. (There are lots of orphaned “Volume I”s about, after all.) Even though the commitment made explicit in the antecedent of (13) *is* the commitment expressed in the consequent, there are claims, such as (11), that are incompatible with its consequent and not incompatible with its antecedent. The difference in content between (9) and (10), which we are accustomed to think of as a difference in truth conditions (compatible with the identity of their assertibility conditions), just is the difference in their consequences, encapsulated in the different status of the conditionals (12) and (13). And that difference manifests itself in a difference in the claims that are incompatible with (9) and (10), a notion we can understand entirely in terms of the normative statuses of commitment and entitlement. Put another way, looking at propositional content in terms of incompatibilities, themselves defined in terms of the fundamental normative statuses of commitment and entitlement, provides the expressive resources to distinguish between the sense of ‘assertible’ that falls short of guaranteeing truth (as ‘foresee’ does), and the sense (perennially sought in terms of some sort of ‘ideal’ entitlement, in a sense of ‘ideal’ that removes it substantially from actual practices of giving and asking for reasons) that would guarantee truth. This is the sense of “It is assertible that...” that would be redundant, in that the incompatibilities associated with “It is assertible that *p*,” would be just those associated with *p*, as they are for “It is true that *p*.”

The point of all this is that the *objectivity* of propositional content—the fact that in claiming that the swatch is red we are not saying anything about who could appropriately assert anything, or about who is committed or entitled to what, are indeed saying

something that could be true even if there had never been rational beings—is a feature we can make intelligible as a structure of the commitments and entitlements that articulate the use of sentences: of the norms, in a broad sense, that govern the practice of asserting, the game of giving and asking for reasons. And we can make sense of practices having that structure even if we understand commitment and entitlement as themselves *social* statuses, instituted by the attitudes of linguistic practitioners. *All* that is required is that the commitments and entitlements they associate with ordinary empirical claims such as “The swatch is red,” generate incompatibilities for these claims that differ suitably from those associated with any claims about who is committed to, entitled to, or in a position to assert anything. The recognition of propositional contents that are objective in this sense is open to any community whose inferentially articulated practices acknowledge the different normative statuses of commitment and entitlement. I argued in the previous section that this includes all *rational* communities—all of those whose practices include the game of giving and asking for reasons. According to the thesis of linguistic rationalism, this is all linguistic communities whatsoever. I have tried here to explain how we can begin to understand the objectivity of our thought—the way in which the contents of our thought go beyond the attitudes of endorsement or entitlement we have toward those contents—as a particular aspect of the normative fine structure of rationality.⁷

⁷ A fuller telling of this story (such as that in *MIE*) would distinguish *three* moves beyond classical assertibility theories, in order to fund a suitable notion of objective representational content for declarative sentences: a) The move from treating assertibility as the fundamental normative pragmatic or force-related notion to commitment and entitlement (which then make it possible to define incompatibility). b) The move from the *circumstances* under which the normative status in question is acquired (=assertibility conditions) to include also *consequences* of acquiring it, as urged in Chapter One. This is moving towards a notion of content as inferential role, identifying propositional contentfulness as suitability to play the role both of conclusion and of premise in inferences of various sorts. The interaction of this move with the previous one generates the three notions of inference (commitment-preserving, entitlement-preserving, and

incompatibility entailments) employed in the text. c) The move from looking at normative *statuses* (assertibility, commitment, entitlement) to normative *social attitudes*. This is to focus on *attributing* (to others) and *acknowledging* (oneself) commitments and so on, as the primary phenomenon. Chapter Five argued that this distinction of *social perspective* is what makes intelligible the specifically *representational* dimension of propositional contents. One might have worried, at the end of that story, about how it is possible (what one has to do in order) to adopt, as it were, a third person perspective towards one's *own* attitudes, and so take them to be subject in principle to the same sort of assessment to which one subjects the attitudes of others, in offering *de re* specifications of their contents. The argument of this chapter provides the answer to that question.