Conceptual Content and Discursive Practice

I

For the past hundred years or more the philosophy of language has been inscribed in a space delineated by two polar approaches. On the one hand, a structural-semantic approach whose avatars include Frege, Carnap, and Tarski focuses primarily on the way the contents expressed by complex or compound expressions[1] depend on those of simpler ones. On the other hand, a pragmatic-anthropological approach characteristic of the American pragmatists, the early Heidegger, and the later Wittgenstein directs its attention in the first instance to the natural history of language as a social practice. The former addresses *meaning*, the latter *use.* Where the first is concerned with what is *said*, the second is concerned with what one is doing in *saying* it. Where one asks what it is that one knows (or believes) when one knows *that* p, the other asks what one must know *how* to *do* in order to count as being in a state that exhibits, or producing a performance that expresses, such a content.

Thought of this way, I take it that the two enterprises should be thought of as complementary rather than competing. For semantics, the theory of the contents

[1] Not the same thing. Compound expressions actually contain the simpler expressions from which they are constructed. Complex expressions, paradigmatically complex predicates (the sort of thing quantifiers attach to), are formed by assimilating compound expressions into equivalence classes of substitutional variants. There is no simple predicate—the sort of thing used to build compounds such as "admired and wrote about"—that "Rousseau admired Rousseau," and "Kant admired Kant" share with each other that they do not also with "Kant admired Rousseau." But the first two share a *complex* predicate "α admired α" that they do not share with "Kant admired Rousseau." Complex expressions are not *parts* of compound ones, but *patterns* compound expressions exhibit.

expressed by using various sorts of vocabulary, and pragmatics, the theory of the practices of using those locutions, can each be pursued in ways that are at most provisionally independent of one another—subject always to the proviso that neither sort of theory can count as adequate unless compatible with an at least acceptable version of the other. Nonetheless, the Fregean and the Wittgensteinean traditions have not as a matter of historical fact had a lot to do with one another. With some honorable exceptions[2], philosophers of language in the model-theoretic tradition of formal semantics do not concern themselves with the issue of what it is linguistic practitioners must be able to *do* in order to associate semantic interpretants of the favored kind—typically construed in rather abstract ways—with the simple expressions of the language. (One excuse for that failure—to my eyes, one that is particularly lame—appeals to the division of intellectual labor: the philosophy of mind is to make up the difference, by offering an account of the thoughts and intentions that accompany speech. But surely the *generic* challenge of semantics encompasses accounting for the contents *both* of thought and of talk.) Conversely, pragmatists and neo-pragmatists, whether of classical, early Heideggerian, or Wittgensteinean lineage, have proven not so much unsuccessful at as simply uninterested in extracting from their investigations of the practices of using linguistic expressions detailed semantic lessons of the sort that abound in, say, possible worlds semantics (concerning for instance, the contents expressed by modal or conditional locutions). Of course the disjunction between semantic and pragmatic concerns is not total. Two of the most discussed approaches in the philosophy of language of our day are those of Davidson and Dummett—philosophers each

[2] I'm thinking of efforts such as that of David Lewis, in "Languages and Language" in K. Gunderson (ed.) *Language, Mind, and Knowledge* [University of Minnesota Press, 1975], and Robert Stalnaker, in *Inquiry* [Bradford Books, 1984].

distinguished precisely by the fundamental and distinctive ways in which they bring semantic and pragmatic concerns to bear on one another. (Outside of the Anglophone tradition, one should in this connection esteem as well the tradition that grew out of Lorenzen's constructivism.)

In *Making It Explicit* I expound an approach to semantics and pragmatics that seeks to do better justice to the insights of the two polar traditions. On the semantic side, it pursues an idea of Frege's: that specifically *conceptual* content is to be distinguished by its relevance to the proprieties of *inference* associated with an expression.[3] The conceptual contents of states, performances, or expressions, should be identified with their inferential roles, that is, their roles in *reasoning*. To be rational is to be sensitive to the normative force of reasons, to tell in practice what is a reason for what, to be able to distinguish good reasons from bad ones. So on this line of thought there is a deep and intimate connection between the concepts <u>concept</u> and <u>rationality</u>, and hence between <u>semantics</u> and <u>reason</u>. According to this *semantic rationalism*, the relation is a reciprocal sense-dependence: one cannot grasp the concept of <u>conceptual content</u> except insofar as one grasps the concept of <u>reasoning</u>, and vice versa.

II

[3] "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." Frege, *Begriffsschrift*, section 3.

How one thinks about *rationality* is then of the first importance for this approach to semantics. The account developed in *Making It Explicit*[4] is best understood against the background of three more familiar models of rationality: the logical, instrumental, and interpretational.

On one picture, to be rational is to be logical. Being sensitive to the force of reasons is a matter of practically distinguishing logically good arguments from those that are not logically good. For a set of claims to serve as a good reason for another claim is for there to be a logically valid argument relating them to that claim as premises to conclusion. Nonlogical facts and the meanings of nonlogical vocabulary contribute to reasoning only by providing premises for logically valid inferences.

The program of assimilating all good reasoning to this model has been immensely influential and productive in the philosophical tradition. It took its modern form when Frege vastly increased the expressive power of logic by giving us formal control over the inferential significance of quantificationally complex properties. The success this idiom was shown to have in codifying mathematical reasoning—by Frege himself, by Hilbert, and by Russell and Whitehead—was a major impetus for logical empiricism, whose central project was to extend the logical model of reasoning to include empirical science. Just when it looked as though the limits of this enterprise had been reached, technical advances in the logical expression of modalities gave the undertaking new life.

[4] Harvard University Press, 1994.

The logical model of reasoning is most at home close to its origins: in codifying *theoretical* inference, the way beliefs can provide reasons for other beliefs. The instrumental model of reasoning begins with *practical* inference—in particular, the way desires or preferences, together with beliefs, can provide reasons for *action*. It identifies rationality with intelligence, in the sense of a generalized capacity for getting what one wants: the reason of Odysseus, rather than of Aristotle. What one has reason to do, on this model, is what provides a means to an endorsed end. Means-end reasoning is formally codified in rational choice theory, in both its decision-theoretic and game-theoretic species. Dutch book arguments show that utility (the measure of preference) will be maximized by practical reasoners who assign probabilities to compound beliefs in ways that satisfy the axioms of classical probability theory. And the laws of classical logic can be deduced as special cases from those axioms. So the instrumental model of rationality has some claim to subsume the logical one as a special case.

One thing to notice about these two models of rationality is that they both treat (nonlogically) contentful beliefs and desires as inputs. *Given* a set of beliefs, and perhaps desires, they purport to tell us which connections among them are *rational*: which constellations of them provide genuine *reasons* for which others. They accordingly presuppose that the *contents* of those psychological states can be made intelligible independently and in advance of considering *rational* connections among them. The idea that one can first fix the meaning or content of premises and conclusions, and only then worry about inferential relations among them, is characteristic of traditional and twentieth century empiricism. This implicit semantic commitment is questioned, however, by the rationalist tradition in semantics, which sees issues of what is a reason for what as essential to the identity and individuation of the conceptual contents that stand in those inferential relations.

The logical and instrumental models of reasons are also (and not coincidentally) alike in their formality. Each sees rationality as being a matter of the *structure* of reasoning, rather than its *content*. The substantial content of the beliefs and desires that provide the premises for candidate theoretical and practical inferences are wholly irrelevant to the rationality of the conclusions drawn from them. *All* that matters for the correctness of the inference is that they have the form of deductively valid inferences or maximization of expected utility *given* those premises. The premises themselves are beyond criticism by these models of rationality, unless and insofar as they themselves were acquired as conclusions of prior inferences, which *are* assessable in virtue of *their* form—and then only relative to the prior (only similarly criticizable) commitments that provide their premises.

A model of rationality that is not in this way purely formal is the translational-interpretational model, most fully developed by Davidson. According to this view, to say that some behavior by others is rational is roughly to say that it can be mapped onto *our* linguistic behavior in ways that make it possible for us to converse with them—at least to draw inferences from their claims, to use them as premises in our own reasoning. The idea is to use our own *practical* know how, our ability to distinguish reasons from nonreasons and to tell what follows from what, to assess the *theoretical* rationality of others. They are rational insofar as their noises (and other behavior, described in nonintentional terms) can be mapped onto ours so as to make them make sense by our standards: to exhibit them as believers in the true and seekers after the good by our own lights. Rationality, then, is by definition what *we*'ve got, and interpretability by us is its definition and measure.

Rationality is not on this view a formal matter at all. For the unintelligibility or wackiness of the substantive, nonlogical beliefs and desires we take our interpretive targets to be evincing in their behavior, both linguistic and nonlinguistic, is every bit as relevant to assessments of their rationality as the connections between them we discern or take them to espouse. We have to be able to count the others as agreeing with us in the contents of and (so) connections among enough of their beliefs and desires to form a background against which local disagreements can be made intelligible, if we are to find them interpretable, that is, rational—for what they have to show up as beliefs and desires—at all.

Rationality as interpretability can also claim to subsume or incorporate both the logical and the instrumental models of rationality. For the first, the explicit form of a Davidsonian interpretation includes a recursive truth theory for the idiom being interpreted, including novel sentential compounds that have never actually been used. So identifying expressions functioning as logical vocabulary can provide a formal framework within which the rest of the interpretive process can take place. Being *logical* creatures is on this view a necessary condition of being *rational* ones, even though there is a lot more to rationality than just that. For the second, making the behavior of the interpreted creatures intelligible requires attributing sample bits of practical reasoning. And Davidson takes it that those will have the form of what he calls "complete reasons": constellations of beliefs and desires that rationalize the behavior according to the instrumental model. Unless one can interpret the target behavior as for the most part instrumentally rational, one cannot interpret it at all.

Finally, the interpretive model does not take the *rational* connections among psychological states or the sentences that express them to be irrelevant to the *contents* they are taken to evince. On the contrary, what makes something have or express the content it does is what makes it interpretable in one way rather than another. And that is a matter of its connections to other things, the role it plays in the overall rational behavioral economy of the one being interpreted. What makes it right to map another's noise onto this sentence of mine, and so to attribute to it the content expressed by that sentence in my mouth, is just that its relations to other noises sufficiently mirrors the relations my sentence stands in to other sentences of mine: what is evidence for and against it, and what it is evidence for and against, as well as what environing stimuli call forth my endorsement of it and what role it plays in practical reasoning leading to nonlinguistic action. Those consequential relations are of the essence of interpretability, and so of rationality on this model.

III

I think the interpretational model of rationality is correct as a *criterion* of rationality, but that it needs to be supplemented in order to yield an *account* of rationality. For a set of practices to be interpretable in this sense means that there is a mapping of it onto *our* discursive practices that preserves the goodness of a whole variety of inferences: theoretical, practical, and reliability inferences (about which more later). Insisting on that is not yet saying what it is for the articulation of *our* practices to count as genuinely *inferential* articulation, and so as capable of conferring *conceptual content* on the states, performances, and expressions that are so articulated and related. *Making It Explicit* aims to offer a *direct* account of a structure such that any set of practices exhibiting that structure thereby counts as *discursive*, and hence as conferring *conceptual content* on things playing suitable roles in those practices. In keeping with its semantic rationalist inspiration, the basic role in question is that of *premise* and *conclusion* in *inferences*. If so, then the primary form of *conceptual* content is *propositional* content, the sort expressed by *declarative sentences*. This is the thought that led Kant to insist that the judgment is the smallest unit of awareness or experience, and Frege to privilege truth values as the most basic sort of semantic intepretant. The rationalist answer to the sometimes vexed question of how to understand "the unity of the proposition" is an inferential one: to be propositionally contentful is to play a distinctive role in a practice of giving and asking for reasons, namely to both be able to serve as and to stand in need of—be offered as and provided with—*reasons*. Playing that dual inferential role is what distinguishes expressions for *propositions* from expressions of other sorts of content.

What, then, are the minimal conditions on being a practice of giving and asking for reasons? It is at this point that the basic *pragmatic* idea of *Making It Explicit* enters the scene. *Discursive* practices are those involving the adoption and attribution of two sorts of *normative* status, *commitments* and *entitlements*. When the practical *consequential* relations among those statuses have the right structure—a matter of the practical consequences of adopting or attributing one for the adoption or attribution of another—they count as *inferentially* articulated. Shifting attention from semantic concern with contents expressed to pragmatic concern with the acts of expressing brings to the fore the performance-kind *asserting* or *judging*. For it is assertings (publicly) or judgings (privately) that can have the pragmatic significance both of being liable to demands for reasons and of serving to respond to such demands. What one is doing in making an assertion or a judgment is undertaking a special kind of *commitment*. (Indeed, as I read him, Kant's reason for taking judgments to be the minimal form of awareness is that they are the smallest unit for which one can be *responsible*.) It is characteristic of that kind of commitment that what one is committed *to* by undertaking an assertional commitment includes *other* such commitments—presystematically, those that follow from it inferentially. Apart from such consequences for other commitments, performing an assertion or judgment would be idle, of no pragmatic significance whatever, a wheel that engaged with no further mechanism. And to say that assertings and judgings are things are liable to demands for reasons is to say that part of what one makes oneself responsible *for* by undertaking such a commitment is, under appropriate circumstances, to vindicate one's *entitlement* to it, by undertaking *other* commitments that serve as reasons for it, i.e. from which it follows.

The basic thought, then, is that for a practice to count as specifically *discursive*, that is as conferring *conceptual* content—which according to the thesis of semantic rationalism means to be a practice of giving and asking for *reasons*—it must involve at least two sorts of normative status, commitments and entitlements, standing in consequential relations to one another. That thought motivates a picture of discursive practice as a kind of *deontic scorekeeping*. To engage in a discursive practice is to keep track of the commitments and entitlements of other practitioners. The pragmatic *significance* of a speech act is the difference it makes to what the performer and his audience are committed and entitled to. One's grasp or understanding of the significance of such a performance is the difference one *takes* it to make to everyone's commitments and entitlements to commitments. So the significance *for* an interlocutor of a speech act or other performance *in* a deontic context consisting of the commitments and entitlements of the various participants is a matter of the way it mandates the *updating* of that context, yielding a subsequent context or 'score'. And the *content* of the expression uttered can be identified roughly with the update *function*, which specifies for each possible deontic context of utterance what the scorekeeping significance of producing that expression (paradigmatically, with assertional force) in that context would be.

The connection between the normative scorekeeping pragmatics and the inferentialist semantics is secured by the idea that the *consequential* scorekeeping relations among expression-repeatables needed to compute the significance updates can be generated by broadly *inferential* relations among those expression-repeatables.[5] The theory propounded in *Making It Explicit* is that there are six consequential relations among commitments and entitlements that are *sufficient* for a practice exhibiting them to qualify as *discursive*, that is, as a practice of giving and asking for *reasons*, hence as conferring *inferentially* articulated, thus genuinely *conceptual* content on the expressions, performances, and statuses that have scorekeeping significances in those practices. The bold and potentially falsifiable overall claim of the whole work is that any practices that exhibit this full six-fold structure will be interpretable in a broadly Davidsonian sense: roughly, mappable onto *ours* in a way that makes *conversation* with us possible. The six sorts of deontically definable relations can be put into two groups: fundamental-semantic and social-pragmatic.

IV

The three fundamental semantic relations are commitment-preserving inferences, entitlement-preserving inferences, and incompatibility entailments. Commitment-preserving inferential relations among the contents expressed by two sentence-repeatables[6] *p* and *q* are imputed by adopting in practice a consequential scorekeeping regularity: that anyone committed to (what is expressed by asserting) *p* is committed to (what is expressed by asserting) *q*. The concept of <u>commitment-preserving inference</u> is a

[5] In *MIE* there are two fundamental sorts of tokening-repeatables: symmetric lexical-syntactic cotypicality equivalence classes and asymmetric anaphoric substitution-inferential inheritance tree structures. See Chapter Seven.
[6] See footnote 4.

generalization of that of *deductive* inference, from the case of *logically* good to the case of *materially* good inferences. Thus any discursive scorekeeper who in practice takes it that as a matter of fact all the yachts in the harbor are sloops will take it that one who claims the *John B.* is in the harbor is thereby committed, whether she knows it or not, to the *John B.* being a sloop. Entitlement-preserving inferential relations among the contents expressed by two sentence-repeatables *p* and *q* are imputed by adopting in practice a corresponding consequential scorekeeping regularity: that anyone *entitled* to a commitment to (what is expressed by asserting) *p* is thereby *prima facie* (more about that later) *entitled* to a commitment to (what is expressed by asserting) *q*. The concept of entitlement-preserving inference is a generalization of that of inductive inference. Thus anyone who endorses the sailor's homily "Red sky at night, sailor's delight; red sky in morning, sailor take warning," will take it that being entitled to the claim that the sky is red at sunset provides a reason for, and in that sense entitles one (other things being equal) to the claim that the weather will be fine tomorrow. It does not yet *commit* one to that claim, for one may also believe that if the barometer falls in the evening, there will be stormy weather the next day, and that the barometer is falling at sunset. Unlike good *de*ductive inferences from true premises, the conclusions of good *in*ductive inferences from true premises may contradict one another.

The third fundamental semantic relation involves the interaction of commitments and entitlements. What is expressed by *p* and *q* may be said to be *incompatible* just in case *commitment* to one precludes *entitlement* to the other. Thus "I will work all day tomorrow in my office," and "I will drive you to the airport tomorrow afternoon," are incompatible in this sense. It is not that I cannot commit myself to both. But if I do, I do not count as *entitled* to either commitment, for the other one in this sense rules it out. Incompatibilities of content generate entailments among them. Thus being a mammal entails being a vertebrate, in that everything incompatible with being a vertebrate (for instance, being a squid) is incompatible with being a mammal, but perhaps not vice versa. (For instance, being a reptile is incompatible with being a mammal, but not with being a vertebrate.) Incompatibility entailments are modally robust: if my favorite pet (an octopus, we may suppose) *were* a mammal, she *would be* a vertebrate.

Here is how these semantic relations among contents help determine the pragmatic significances of (paradigmatically assertional) speech acts involving expressions with those contents. When *S* asserts *p*, scorekeepers first add to the repertoire of commitments attributed to *S* commitments to all those claims *q* such that (according to the scorekeeper) there is a commitment-preserving inference from *p* to *q*. (This will obviously include *p*.) Next, one adds to the repertoire of entitlements attributed to S all the *q*'s such that (according to the scorekeeper) there is an entitlement-preserving inference from *p* to *q*. Finally, one *subtracts* from the repertoire of entitlements attributed to S all the claims *q* such that there is some *r* incompatible with *q* to which *S* is (according to the scorekeeper) committed.

This last feature is the primary locus of the important difference between the sort of (social) *normative* functionalism about semantic content put forward in *Making It Explicit* and more familiar (individual) *causal* functionalisms. For accounts of content in terms of causal roles have difficulty making sense of the possibility of inconsistent, or more generally, incompatible beliefs. One cannot, after all, simultaneously be in states individuated in part by the fact that they cause incompatible behaviors. Such theories sometimes, implausibly, simply deny the possibility of having inconsistent beliefs. (Possible comparison: the denial by rational choice theorists that it could be so much as intelligible to interpret behavior by attributing cyclical preferences.) Otherwise, causal-functional theories typically make *ad hoc* adjustments, either by partitioning believers with inconsistent beliefs into competing consistent subsystems, or to merely statistical causal influences. By contrast, no corresponding surds threaten the normative functionalist conception of agents with incompatible *commitments*. The possibility of undertaking (and hence of attributing) incompatible doxastic commitments is straightforwardly intelligible in the same way that undertaking incompatible practical commitments is. Just as I can, if I am sufficiently thoughtless, injudicious, or just unlucky, make incompatible promises concerning what I will do, and so how things are to be—at the cost, to be sure, of being unable to fulfill them all—so I can make incompatible commitments concerning how things are. The corresponding cost is that I cannot then be *entitled* to any of those commitments. But that fact does not by itself in any way undercut my status as nonetheless committed.

One of the principal ways in which I can non-culpably acquire incompatible commitments is when some of the inferential consequences of prior commitments to which I am entitled collide with the current deliverances of my non-inferential reporting capacities. Consider a microbiologist who has concluded, on the basis of symptoms and exposure, that the rod-shaped bacteria recovered from a diseased tissue is *E. coli*, a Gram-negative bacillus. If she then finds that it turns purple when exposed to crystal violet stain, and hence is Gram-positive, she has acquired incompatible beliefs, one as the result of inference, and the other non-inferentially, by inspection of the slide. Response to the stain indicates *B. cereus*, but the presence of that bacillus is, we may suppose, incompatible with the observed symptoms and demonstrated exposure. Semantics dictates only that she cannot be *entitled* either to the claim that the agent is *E. coli* or to the claim that it rather *B. cereus*. It does not tell her what to do to repair this situation of incompatible commitments—for instance, whether to suspect that the symptoms were misdescribed, the exposure misreported, or the stain misapplied. The details of that sort of updating are a cognitive, not merely a semantic matter.[7]

V

This example indicates a crucial mechanism whereby features of the causal order can come to be reflected in the conceptual order. The essential *empirical* dimension of conceptual content depends not only on the fundamental semantic notion of <u>incompatibility</u> (definable, as we have seen, in terms of the still more basic normative status concepts <u>commitment</u> and <u>entitlement</u>), but also on the distinctive semantic contribution of non-inferential reports. That topic brings us to the second triad of

[7] And this is true even if the incompatibility of commitments is more diffuse. For material incompatibilities can be like formal logical ones, in that one may have a set of n commitments, any n-1 of which are compatible, but which jointly are incompatible—as for instance is true on the logical side of {p, p→q, q→r, ~r}. Semantically, anyone with all these commitments is entitled to *none* of them. But that semantic fact only settles it that one ought to do *something* to repair the situation, undertake *some* revision of those commitments. It does not offer advice about which of the many ways of doing that might be best.

deontically definable content-constituting relations: the social-pragmatic ones. The three basic social pragmatic structures are the *empirical*, the *practical*, and the *testimonial*. The fundamental semantic relations were all strictly inferential, and reflecting practically what Sellars calls "language-language moves".[8] The social pragmatic relations widen the practical focus. The *empirical* dimension articulates the contribution to conceptual content made by language-*entry* moves, in which a knower responds to a (typically, but not always) *non*-linguistic situation by acknowledging an inferentially significant doxastic commitment. The *practical* dimension articulates the contribution to conceptual content made by language-*exit* moves, in which a knower responds to the acknowledgment of an inferentially significant practical commitment by altering a (typically, but not always) non-linguistic situation. The first encompasses non-inferentially elicited perceptual reports, and the second non-inferentially efficacious intentional agency. The *testimonial* dimension articulates the contribution to conceptual content made by the possibility of *interpersonal* non-inferential inheritance of commitments and entitlements.

Our capacity to find out how things are in the world around us depends on our ability to practically distinguish different kinds of environing states of affairs by responding differentially to those stimuli. In the simplest cases, there is nothing distinctively *cognitive* or *discursive* about such abilities. Chunks of iron rust in wet environments and not in dry ones, branches of a given diameter break with sufficiently heavy loads and not with lighter ones, motion detectors turn on the porch light if large things move across their field of sensitivity, but not if small ones do, and so on. By reliably responding

[8] In "Some Reflections on Language Games" [ref.].

differentially to different kinds of stimuli, such systems can be understood as *classifying* their environments as being of different kinds. But this is not yet *conceptual* classification, and so not yet conceptual *awareness*. The boundary from mere sensitivity, irritability, or even sentience, to true sapience is crossed when the response that is reliably differentially keyed to stimuli of a certain kind is the application of a *concept*—the use of a word that plays an appropriate role in *inferential* relations of the three sorts already discussed. The difference between a parrot trained to respond by uttering the noise "That's red!" to the presence of visible red things, on the one hand, and a genuine reporter of red things, on the other, is that reporter's utterance is, as the mere responder's is not, the undertaking of an inferentially articulated *commitment*: the making of a move in a game of giving and asking for reasons, the endorsement of a content that can both serve as a reason for endorsing others and is liable to demands for vindication of the reporter's entitlement to a commitment with that content. Empiricists have always been right to stress that apart from the capacity reliably to respond differentially to stimuli in perception, not only empirical knowledge, but even empirical conceptual content would be unintelligible. And rationalists have always been right to stress that apart from the role a performance, state, or expression plays in reasoning, its conceptual contentfulness is unintelligible. As Kant taught us, they err when they succumb to their characteristic complementary temptations to treat one or the other of these necessary conditions of content, and so cognition, as sufficient. The causal dimension of reliable responsiveness and the inferential dimension of rational endorsement are equally essential to empirical conceptual contentfulness.

The two aspects of empirical concepts are to a certain extent intelligible independently of one another. We've just seen that reliable differential responsive dispositions can be discerned in non-discursive systems. And the furthest reaches of mathematics—abstract algebra or pure set theory, for instance—give us a grip on inferential relations that in principle need not be anchored in perceptual or practical interactions with concrete things. Nonetheless, the two aspects—causally reliable responsiveness and inferentially articulated commitment—should not be understood as simply bolted together to yield non-inferential reporting capacities. For they interact intimately. The basic idea is that at least some inferential relations must be brought to track causal ones, so that reliable responsiveness is incorporated in the inferential articulation of conceptual contents.

To see how this works, think about the semantics in terms of the pragmatics (inferential relations in terms of scorekeeping activities). What must one *do* in order to be taking or treating another interlocutor as a *reliable* reporter, of, say, red things? To take someone to be a reliable reporter is taking it that he is likely to be right, that his observational claim is likely to be *true*. But taking a claim to be true is just endorsing it oneself. That is, to treat the other as reliable is to take it that his non-inferentially elicited response to something *as* red—his endorsement of the observation report—provides a *reason* for endorsing that claim oneself. To do that is to endorse an *inference* of a distinctive kind, a *reliability* inference. That is an inference from an *attributed* commitment to a corresponding *acknowledged* commitment. More carefully put, for S to take S' to be a reliable reporter of red things is for S to keep deontic score in such a way that whenever S takes S' to be responsively *committed* to the claim that something is red, everyone else, S included, is thereby counted as *prima facie entitled* to that same claim. The entitlement in question is only *prima facie,* since any entitlement is defeasible by concomitant commitment to incompatible claims. And for the same reason, S need not always endorse the claim attributed to S' in order to treat S' as reliable, since in some cases S may have collateral commitments incompatible with that one. Thus I may take you to be a generally reliable reporter of red things and still not endorse your report of a particular object as red, if I have information you lack about the non-standard lighting conditions under which you are viewing it. For your authority to be *genuine*, it need not be *indefeasible*.

The reliability of the *causal* process by which a reporter responds to the *fact* of visible red things by endorsing the *claim* that there are red things there is then reflected in *conceptual*, that is to say, *inferential* relations by the consequential scorekeeping relation between *attributing* commitments to the reporter and *acknowledging* corresponding entitlements for the reliability-attributing scorekeeper and (according to that same scorekeeper) also for others. The reliability inference is an *inference*, but unlike committive, permissive, and incompatibility entailments, it is an essentially *socially* articulated inference. For it relates commitments *attributed* to a reporter to entitlements *acknowledged* by the scorekeeper, and attributed not only to the reporter, but also to others in the reporter's audience.

This social articulation of the reliability inference shows up the authority of reliable observers as a species of a more general kind of interpersonal entitlement-inheritance:

testimony. The dimension of testimonial authority is an essential element of the deontic scorekeeping understanding of the pragmatic significance of the fundamental (because pathognomic for the characterization of a practice as specifically *discursive*) speech act of assertion. The basic thought is that in asserting that *p*, one is doing two things. On the side of *authority*, one is *licensing* others to reassert one's claim. On the side of *responsibility*, one is committing oneself to justify the claim if suitably challenged. Both of these should be understood in terms of the inheritance of *entitlement* to the commitment one has undertaken. A scorekeeper who acknowledges the authority of one's assertion will count those in the audience as having inherited from it *prima facie* entitlement (defeasible, as always, by concomitant commitment to incompatible contents) to the content one has endorsed. And if one is challenged, paradigmatically by the assertion of an incompatible claim to which the challenger is at least *prima facie* entitled, one is then obliged, in the eyes of the scorekeeper, to vindicate one's entitlement to the commitment undertaken in the original assertion, paradigmatically by undertaking other commitments that stand to it as premise to conclusion in an at least entitlement-preserving inferential relation.

What if one fails to fulfill that justificatory responsibility in response to a suitably qualified challenge? The cost of that failure is simply the loss of the authority one implicitly claimed by issuing the assertion in the first place. For a scorekeeper who takes it that one has not shown oneself to be entitled to the commitment in the first place thereby takes it that one does not have a heritable entitlement to bequeath to others. On the other hand, while that re-assertion license *is* in effect, either because it has not been challenged or because any actual challenges have been satisfactorily responded to (all in the eyes of the scorekeeper to whom we are attending), the force of the authorizing is that those who by its means inherit *prima facie* entitlement to the claim can, if *they* are challenged, vindicate their entitlement by *deferral* to the original assertor, in lieu of having to produce independent justifications. Thus on this picture the interpersonal, intracontent inference-licensing authority of an assertion, and the intrapersonal, intercontent inferential-justificatory responsibility one undertakes by implicitly asserting that authority are co-ordinate and interdependent: two sides of one coin.

We saw that the authority of reliable observers is presumptive, but defeasible. So too with the more general category of testimonial authority. The picture is that assertions by competent speakers have a default authority. Their normative power to license or entitle others to reassert the claim, to use it as a premise in their own inferences, is considered to be in force until and unless their assertion is challenged by an incompatible claim to which the challenger has at least an equivalent *prima facie* entitlement. (That would be lacking if the challenger were also committed to claims incompatible to the challenging one, for instance.) For granting that default *prima facie* entitlement to assertions just on the basis of the fact that the commitment has been undertaken is what it is for a scorekeeper to treat the assertor as a competent deployer of the concepts in question—whether they are observation concepts or not.

These observations about the contribution of pragmatics to semantics in the system under discussion make it possible to show something about the relation between the various sorts of fundamental inferential relations that is not otherwise obvious. In particular, it is possible to show that if one claim incompatibility entails another, then there is also a good commitment-preserving inference from the first to the second. Suppose that *p* incompatibility entails *q*. This means that everything incompatible with *q* is incompatible with *p*, and suppose that *S* is committed to *p*. The question is whether *S* thereby counts as committed to *q*. We can start by asking what it means in scorekeeping terms to be committed to *q*. We've seen that that status is specified in terms of the two dimensions of authority and responsibility it involves: First, undertaking a commitment to a claim involves *authorizing* or licensing others to assert it (undertake commitment to it), which in practice means authorizing them to defer to the original assertor the responsibility to justify or demonstrate entitlement to the claim if appropriately challenged. Second, (and presupposed by the first), undertaking commitment to a claim involves the *responsibility* to justify or demonstrate entitlement to it if it is appropriately challenged. From the point of view of our current concerns, the important thing to notice about this articulation of the pragmatic scorekeeping significance of doxastic *commitments* is the crucial role played in it by appeal to *entitlements*. It is this, ultimately, that connects incompatibility entailments with commitment-preserving inference (and indeed, the latter with entitlement-preserving inference).

With this reminder on board, consider whether commitment to *p* involves commitment to *q* wherever *p* incompatibility entails *q*. Think first of the dimension of responsibility: If one is committed to justify *p* when it is appropriately challenged, is one thereby committed to justify *q* when it is appropriately challenged? An appropriate challenge is the (at least *prima facie* justified) assertion of something incompatible. But since by hypothesis *p* incompatibility entails *q*, everything incompatible with *q* is incompatible with *p*. So every challenge to *q is* a challenge to *p*. So one must be able to respond to all these challenges in order to respond to all the challenges to *p*. That is, in order to be able to justify *p*, one must be able to justify *q*. So on the side of responsibility, commitment to *p* involves commitment to *q*. What about the dimension of authority? There too, in authorizing others to assert *p*, one has thereby authorized them to assert *q*. For the cash value of that authorization is authorization to defer justificatory responsibility, in response to warranted challenges. And again, by hypothesis any warranted challenge to *q* is a warranted challenge to *p*. So in taking on the responsibility to answer for *p*, one has thereby also taken on the authority to answer for (respond to challenges to) *q*. Thus along both the dimension of authority and that of responsibility, commitment to *p* involves commitment to *q*, if *p* incompatibility entails *q*.

We can see in much the same way why if there is a commitment-preserving inference from *p* to *q*, there is also a (*prima facie*) entitlement-preserving inference from *p* to *q*. For if everyone who is committed to *p* is thereby committed to *q*, then on the side of authority, in authorizing others to assert *p*, I am thereby authorizing them to assert *q*. And to say that is to say that they can inherit *entitlement* to *q* from my entitlement to *p*. And on the side of responsibility, in undertaking the responsibility to justify or other wise vindicate my entitlement to *p*, I am thereby undertaking the responsibility to justify, or otherwise vindicate my entitlement to *q*. So I cannot be entitled to *p* unless I am also entitled to *q*, which is to say that there is a good inference from *p* to *q* preserving *prima facie* entitlements.

So the social pragmatics of testimonial entitlement inheritance settles it that the three flavors of fundamental semantic inference are strictly ordered in strength. If *p* incompatibility entails *q* then there is a good commitment-preserving inference from *p* to *q*. And if there is a good commitment-preserving inference from *p* to *q*, then there is a good entitlement-preserving inference from *p* to *q*. This is another crucial way in which the normative pragmatic account of what one is *doing* in *committing* oneself to a claim is inextricably bound up with the inferential semantic account of the conceptual contents to which one thereby becomes committed.

VI

The third and final sort of social pragmatic inferential relation is that of *practical* inference. By that I mean inferences whose conclusions include *practical* commitments: commitments to *do* something. These are commitments to *make* some claim true, rather than commitments whereby one *takes* some claim to be true. In the normative pragmatics of *Making It Explicit*, the concept of <u>doxastic commitment</u>, picking out the kind of thing expressed by assertions, does much of the explanatory work for which the concept of the intentional state kind *belief* is called into service by more traditional philosophies of mind. In much the same way, the concept of <u>practical commitment</u> does much of the explanatory work for which the concept of the intentional state kind *intention* is called into service by those traditional theories.

The idea is that being a competent agent is being able reliably to respond to acknowledging at least the most basic sort of practical commitment by bringing about a state of affairs of the sort specified by its content: making true a claim with that prepositional content. This capacity for reliable language-exits is to be understood by strict analogy to the capacity for reliable language-entries. We can be trained reliably to respond to visible red things by acknowledging a commitment to their being red, and because of the way the world is and the way we are wired up we can*not* be so trained non-inferentially to report without the aid of instruments the presence of things that emit radio waves or that once belonged to the emperor. Just so we can be trained reliably to respond to a commitment to raise our arms by making it true that our arms go up, and because of the way the world is and the way we are wired up can*not* be so trained non-inferentially to bring it about without the aid of instruments that we emit radio waves or that the Moon turns a different face toward the Earth. As in the theoretical case of observation, attributing practical reliability to an agent with respect to a range of possible performances is a matter of the scorekeeper adopting a certain sort of inferential commitment. In both cases, the inference is one in which a commitment *attributed* to another is taken as reason for *endorsing* a claim oneself. In this case, my taking you to be reliable arm-raiser is taking it that *your* undertaking a practical commitment to raise your arm gives *me* good reason (entitles me) to commit myself to the claim that your arm will go up.

Practical inferential relations then can be thought of as governing transitions (commitment or entitlement inheritance) from doxastic to practical commitments, that is, from the commitments acknowledged in assertions to commitments to *do* something, made explicit by saying something like "I shall φ." Practical reasoning of this sort is what is rehearsed in deliberation, and attributed in assessments of what others had good reason to do. We make each other's behavior rationally intelligible by attributing and assessing sample bits of practical reasoning that would rationalize what they do, even in the cases where we do not take it that the behavior in question was the result of explicit deliberation, that is, causally resulted by detachment of a practical conclusion from such a process of practical reasoning.

Seen from this angle, expressions of preference or desire show up as codifying commitment to the propriety of patterns of practical inference. Thus S's preference or desire to stay dry is a commitment to the propriety (here, in the sense of entitlement-preservation) of inferences of the form:

Only doing A will keep me dry.

∴ I shall do A.

Attributing a preference or desire is attributing commitment to such a pattern of practical inferences. The explicit claim "I prefer to stay dry," stands to such implicit commitments to patterns of practical inference in the same expressive relation in which the conditional $p \rightarrow q$ stands to a commitment to the propriety of theoretical inferences from *p* to *q* (with different flavors of inference corresponding to different conditionals). In both cases, it is a mistake to confuse the statements that make inference licenses explicit with *premises* required for the inference to be licit in the first place—for reasons Lewis Carroll has made familiar in “Achilles and the Tortoise.”

Further, preferences and desires are only *one* sort of practical inference license. For in general, this is the expressive role distinctive of *normative* vocabulary as such. Thus a statement of the obligations associated with some institutional status, such as “Civil servants are obliged to treat the public with respect,” licenses inferences of the form:

Doing A would not be treating the public with respect.

∴ I shall not do A.

This institutional pattern of practical inference differs from the preference pattern in that the latter is binding only on those who endorse the preference in question, while the latter is binding on anyone who occupies the status in question, i.e. on civil servants—regardless of their desires. Another pattern of practical reasoning is codified by normative claims that are not conditioned upon occupation of an institutional status. Thus “It is wrong to (one ought not) cause pain to no purpose,” licenses inferences of the form:

Doing A would cause pain to no purpose.

∴ I shall not do A.

Endorsing the unconditional normative claim is committing oneself to the bindingness of this form of practical inference for *anyone*, regardless of their preferences or institutional status.

On the inferentialist picture, *all* of these 'oughts'—the instrumental, the institutional, and the unconditional—are in the most basic sense *rational* oughts. For they codify commitments to patterns of practical *reasoning*. From this point of view, the humean, who insists on assimilating all practical reasoning to the first, or instrumental model, on pain of a verdict of practical irrationality, and the kantian, who insists on assimilating all practical reasoning to the third, or unconditional model, on pain of a verdict of practical irrationality in the form of heteronomy, are alike in pursuing Procrustean explanatory strategies. The real questions concern the justification of normative commitments of these various forms: the circumstances under which one or another should be endorsed, and what considerations speak for resolving incompatibilities among such commitments in one way rather than another.

Some vocabulary plays a distinctive role in explicitly marking a commitment as the result of a language-entry move: "I see that the light is red." Similarly, some vocabulary plays a distinctive role in explicitly marking a commitment as the origin of a language-exit move: "I shall raise my arm." Some vocabulary has an immediate observational role: 'red', 'square', and 'cat', but not (usually) or 'emits radio waves' or 'owned by the emperor'. Some vocabulary has an immediate practical role: 'painful', 'cruel', and 'desireable', but not (usually) 'rapid' or 'distant'. But the conceptual contents of a great deal of vocabulary that does not have this sort of special status are nonetheless affected by their inferential links, typically through seriously multipremise inferences, to vocabulary that does. So the content even of concepts at some inferential distance from reasons provided by observation and reasons for action nonetheless has an empirical and practical aspect, in virtue of those inferential connections.

VII

I have now sketched six kinds of consequential scorekeeping relations among commitments and entitlements. The three fundamental semantic ones are consequential commitment, consequential (*prima facie*) entitlement, and incompatibility entailment. We've seen that these are strictly ordered by strength: if *p* incompatibility entails *q*, then anyone who is committed to *p* is committed thereby to *q*, and if anyone who is committed to *p* is committed thereby to *q*, then anyone who is provisionally (that is, apart from incompatible collateral concomitant commitments) entitled to a commitment to *p* is provisionally entitled to a commitment to *q*. The three social pragmatic consequential relations are language-entries through observation, language-exits through intentional action, and the testimonial structure of authority and responsibility that defines the basic pragmatic significance of assertional speech acts—thereby defining the concepts both of <u>propositional content</u>, on the one hand, and <u>declarative sentence</u> on the other. Each of these is to be understood in scorekeeping terms by appeal to *reliability-inferential* commitments on the part of scorekeepers. Inferences of this sort are essentially *social* or *interpersonal* in that they related normative statuses attributed to one individual to those attributed to others or undertaken by the scorekeeper herself. Inferences of this sort are also normatively *intermodal* in that they relate *commitments* to *entitlements*.

The core of the theory of discursive practice and conceptual content put forward in *Making It Explicit* is that these six consequential relations among commitments and entitlements are *sufficient* for a practice exhibiting them to qualify as *discursive*, that is, as a practice of giving and asking for *reasons*, hence as conferring *inferentially* articulated, and so genuinely *conceptual* content on the expressions, performances, and statuses that have the right kind of scorekeeping significances in those practices. This theory is developed to begin with from the point of view of someone looking at a set of social practices from the outside, and asking the question: what must be true of those practices, what structure must they be taken to exhibit, so that understanding them that way *is* implicitly taking them to be *discursive* practices? By the end of the book, however, this sort of external interpretive stance—what one must do, how one must treat an alien community in order thereby to count as taking them to be making *assertions* and *inferences*—is seen to be equivalent to an internal scorekeeping stance *within* a discursive community. That is, one must adopt toward the practitioners the same sort of attitude one both takes them to adopt towards each other and adopts towards one's own discursive fellows. One must keep deontic score on them, attributing commitments and entitlements that both have authority for one's own, and to which one's own commitments and entitlements are responsible. This collapse of levels can be thought of indifferently as the interpreter entering into the community to whom discursive practices are being attributed and as the interpreter treating those practitioners as members of her own discursive community. And that is to say that the interpreter, by becoming a scorekeeper, enters into a dialogical relation with the community being interpreted, mapping their utterances onto her own in a way whose adequacy is to be assessed by the fluency of conversation it enables. In short, the stance in question is a translational-interpretive stance that evidently belongs in a box with the orthodox Davidsonian variety.

The overall claim of *Making It Explicit*, then, is that the six broadly inferential structures outlined in the body of this paper articulate the fine structure of rationality in the Davidsonian interpretational sense. More specifically, the claim is that exhibiting the six-fold structure relating discursive scorekeeping practices to conceptual contents construed as inferential roles is necessary and sufficient for being interpretable, and hence rational, in the Davidsonian sense. That structure purports to be the structure of discursive practice as such—the structure distinctive of practices that deserve to count as practices of giving and asking for *reasons*, and hence (according to the semantic rationalism that is at the core of the theory) as conferring genuinely *conceptual* content. If it is correct, that is the structure that communities of extraterrestrials or of digital computers, however unlike us they may be in other respects, must exhibit if they are to qualify as potential interlocutors—as sapient knowers and agents, ensdorsers of claims and aims, makers and takers of reasons, seekers and speakers of truths.

Bob Brandom