

6/20/2011

Chapter in *The Cambridge Handbook of Linguistic Anthropology*

**Intentionality and Language:
A Normative, Pragmatist, Inferentialist Approach**

Robert B. Brandom

I. Intentionality

In this essay I present a battery of concepts, distinctions, terminology, and questions that are common currency among philosophers of mind and language who think about intentionality. Together, they define a space of possible explanatory priorities and strategies. In addition, I sketch a systematic, interlocking set of commitments regarding the relations among these concepts and distinctions, which underwrites a distinctive set of answers to some of the most important of those questions. This normative, pragmatist, inferentialist approach to intentionality and language is much more controversial. I have developed and expounded it in a number of books over the past two decades. In the present context its exposition can serve at least to illustrate how one might assemble a framework within which to think about the relations among these important issues.

The contemporary philosophical use of the medieval scholastic term “intentionality” was introduced by Franz Brentano [1838-1917]. His student Edmund Husserl [1859-1938] recognized it as apt to characterize a phenomenon that Immanuel Kant [1724-1804] had put at the center of our thought about mindedness, as part of what we would now call his semantic transformation of René Descartes’s [1596-1650] epistemological turn in the philosophy of mind. This is the idea of a kind of *contentfulness* that is distinctive of at least some of our psychological states and linguistic utterances. Brentano characterized intentionality in terms of “reference to a

content, a direction upon an object.”¹ John Searle [b. 1932] offers this pre-theoretical summary of the subject-matter of his book *Intentionality*:

...if a state S is Intentional then there must be an answer to such questions as:

What is S about? What is S of? What is it an S that?²

We can specify the content of someone’s belief by saying, for instance, that she believes *that* Kant’s servant was named ‘Lampf’. In that case, it is a belief *of* or *about* Kant’s servant, *representing* him as being so-named. Brentano was impressed by the thought that while things can only stand in physical or causal relations to actually existing facts, events, and objects, intentional states can “refer to contents” that are not true (do not express actual facts) and be “directed upon objects” that do not exist.³ I can only kick the can if it exists, but I can think about unicorns even if they do not.

We should distinguish *intentionality* in this sense from *consciousness*. These phenomena only overlap. For, on the one hand, pain is a paradigmatically conscious phenomenon. But pains are not in the sense relevant to intentionality *contentful* states or episodes. They do not have contents that could be expressed by sentential ‘that’ clauses. And they are not (at least not always) *about* anything. On the other hand, there is nothing incoherent about the concept of *unconscious* beliefs—which do have intentional contents specifiable both in terms of ‘that’ and ‘of’. Attributions of belief answer to two kinds of norms of evidence, which in some cases diverge. Evidence derived from sincere avowals by the believer license the attributions of beliefs of which the believer is conscious. But beliefs, desires, and other intentional states can also be attributed on the basis of what relatively stable beliefs and desires provide premises for bits of practical reasoning that make the most sense of what the believer actually does, even in the absence of dispositions sincerely to avow the intentional states in question. Where such intentional explanations are good explanations, the attributed intentional states are unconscious.

The need to make this distinction is a manifestation of a deeper distinction between two sorts of mindedness: sentience and sapience. Sentience is awareness in the sense of being

¹ Franz Brentano, “Psychology from the Empirical Standpoint,” trans. D.B.Terrell, quoted on pp. 119-20 in H. Morick (ed.) *Introduction to the Philosophy of Mind: Readings from Descartes to Strawson* [Scott, Foresman; Glenview, Ill. 1970].

² John Searle, *Intentionality* [Cambridge University Press, 1983].

³ Notice that it is at least not obvious that the first part of this claim is true. Reinforcing the dam might have averted a possible disaster. If so, the nonexistence of the disaster was presumably an effect caused by the reinforcement.

awake. Anything that can feel pain is sentient. Sapience is having intentionally contentful states such as beliefs, desires, and intentions: believing, desiring, or intending *of* the dog *that* it is sitting, will sit, or should sit. An essential element of Descartes's invention of a distinctively modern conception of the mind was his assimilation of *sensations* (for instance, pain) and *thoughts* (for instance, that foxes are nocturnal omnivores). His predecessors had not been tempted by such an assimilation of sentience and sapience. His innovation, and the rationale for the assimilation, was an *epistemic* criterion of demarcation of the mental. Both sensations and thoughts, he took it, were *transparent* and *incorrigible* to their subject: they could not occur without the subject knowing that they occurred, and if the subject took it that they occurred, then they did. Apart from growing appreciation (beginning already with Gottfried Leibniz [1646-1714]) of the potential explanatory significance of unconscious mental states, concerning which subjects do not have the sort of privileged epistemic access Descartes focused on, we have come to appreciate the importance of not prejudging issues concerning the relations between sentience and sapience. In particular, we have come to see that some of the most important issues concerning the plausibility, and even the intelligibility, of artificial intelligence as classically conceived, turn on the question of whether sapience presupposes sentience (which is, as far as our understanding so far reaches, an exclusively biological phenomenon).

II. Representational and Propositional Dimensions of Practical and Discursive Intentionality

Within the general area marked out by the term 'intentionality', there are two distinctions it is important to keep in mind: the distinction between practical and discursive intentionality, and the distinction between propositional and representational intentionality. Practical intentionality is the sort of directedness at objects that animals exhibit when they deal skillfully with their world: the way a predator is directed at the prey it stalks, or the prey at the predator it flees. It is a phenomenon of sentience, with the role objects, events, and situations play in the lived life of an animal providing the practical significances (food, threat...) that can be perceptually afforded. At the most abstract level of description, however, biological practical intentionality is an

instance of a kind of broadly teleological directedness at objects that also has non-sentient examples. For any process that has a Test-Operate-Test-Exit feedback-loop structure, where operations on an object are controlled by information about the results of previous operations on it that are repeated until a standard is satisfied, can be seen as in a distinctive way “directed at” the objects the system both operates on and is informed about. This genus includes both finite-state automata executing conditional branched-schedule algorithms, for instance, in a radar-guided tracking anti-aircraft missile, and the fly-wheel governors that regulated the boiler-pressure of the earliest steam engines. Discursive intentionality is that exhibited by concept-users in the richest sense: those that can make judgments or claims that are *about* objects in the semantic sense. The paradigm of the sort of sapience I am calling “discursive intentionality” is exhibited by language users: ones who can *say* what they are thinking and talking about.

The distinction between representational and propositional intentionality is that between two dimensions of content intentional states can exhibit, corresponding to two of Searle’s questions, quoted above: “What is S of? What is it an S that?”. The answer to the first sort of question is the specification of an object represented by the state (“It is a belief of or about ships, shoes, sealing-wax...”), while the answer to the second sort of question is the specification of what is believed or thought (“It is the belief that ships should be sea-worthy, that shoes are useful, that sealing-wax is archaic...”). The first expresses what we are thinking or talking *about*, and the second what we are thinking or saying (about it).

This distinction of two dimensions of contentfulness applies both to the practical and to the discursive species of intentionality. The dog believes *that* his master is home, and he believes *that of* Ben, his master. The principled difficulties we have with using the terms appropriate to discursive intentionality to specify precisely the propositional contents exhibited in practical intentionality (the dog does not really have the concepts specified by “master” and “home”—since it does not grasp most of the contrasts and implications essential to those concepts) do not belie the fact there is some content to his beliefs about that human, Ben, in virtue of which his belief that his master is about to feed him differs from his belief that his master is home, or that someone else will feed him.

Two opposed orders of explanation concerning the relations between practical and discursive intentionality are pragmatism and platonism. Pragmatism is the view that discursive

intentionality is a species of practical intentionality: that knowing-that (things are thus-and-so) is a kind of knowing-how (to do something). What is explicit in the form of a principle is intelligible only against a background of implicit practices. The converse order of explanation, which dominated philosophy until the nineteenth century, is a kind of intellectualism that sees every implicit cognitive skill or propriety of practice as underwritten by a rule or principle: something that is or could be made discursively explicit. A contemporary version of platonism is endorsed by the program of symbolic artificial intelligence, which seeks to account for discursive intentionality as a matter of manipulating symbols according to definite rules. A contemporary version of pragmatism is endorsed by the program of pragmatic artificial intelligence, which seeks to account for discursive intentionality by finding a set of nondiscursive practices (practices each of which can be exhibited already by systems displaying only practical intentionality) that can be algorithmically elaborated into autonomous discursive practices.⁴ Pragmatism need not take the reductive form of pragmatic AI, however.

What about the explanatory priority of the representational and propositional dimensions of intentionality? Here, too, various strategies are available. My own approach is to give different answers depending on whether we are talking about practical or discursive intentionality. Within practical intentionality, the propositional dimension should be understood in terms of the representational dimension. Within discursive intentionality, the representational dimension should be understood in terms of the propositional. (Notice that the possibility of such a view would not even be visible to a theorist who did not make the distinctions with which I began this section.) The sort of representation that matters for understanding practical intentionality is the mapping relation that skillful dealings produce and promote between items in the environment and states of the organism. The usefulness of map representations depends on the goodness of inferences from map-facts (there is a blue wavy line between two dots here) to terrain-facts (there is a river between these two cities). The propositional content of the map-facts is built up out of representational relations that are sub-propositional (correlating blue lines and rivers, dots and cities). Such relations underwrite the representation-to-proposition order of explanation at the level of practical intentionality.

⁴ I discuss these programs in more detail in Chapter Three of *Between Saying and Doing: Towards an Analytic Pragmatism* [Oxford University Press, 2008].

The considerations that speak for this order of explanation for *practical* intentionality are sometimes thought to speak for the same order of explanation for *discursive* intentionality. And the case could only get stronger when one conjoins that commitment with a pragmatist order of explanation relating practical and discursive intentionality. Nonetheless, I think there are strong reasons to endorse the explanatory priority of the propositional to the representational dimensions of intentionality at the level of discursive intentionality. They derive to begin with from consideration of the essentially *normative* character of discursive intentionality.

III. The Normativity of Discursive Intentionality

Kant initiated a revolution in thought about discursive intentionality. His most fundamental idea is that judgments and intentional doings are distinguished from the responses of nondiscursive creatures in that they are things the subject is in a distinctive way *responsible* for. They express *commitments*, or *endorsements*, they are exercises of the *authority* of the subject. Responsibility, commitment, endorsement, authority—these are all *normative* concepts. In undertaking a theoretical or practical discursive commitment that things are or shall be thus-and-so, the knower/agent binds herself by rules (which Kant calls “concepts”) that determine *what* she thereby becomes responsible for. For instance, in making the judgment that the coin is copper, the content of the concept copper that the subject applies determines that she is committed (whether she knows it or not) to the coin’s conducting electricity, and melting at 1085° C., and that she is precluded from entitlement to the claim that it is less dense than water. The difference between discursive and nondiscursive creatures is not, as Descartes had thought, an *ontological* one (the presence or absence of some unique and spooky sort of mind-stuff), but a *deontological*, that is, normative one: the ability to bind oneself by concepts, which are understood as a kind of *rule*. Where the pre-Kantian tradition had focused on our grip on concepts (is it clear, distinct, adequate?), Kant focuses on their grip on us (what must one do to subject oneself to a concept in the form of a rule?). He understands discursive creatures as ones who live, and move, and have their being in a *normative* space.

The tradition Kant inherited pursued a bottom-up order of semantic (they said “logical”) explanation that began with concepts, particular and general, representing objects and properties. At the next level, they considered how these representations could be combined to produce propositions of different forms (“Socrates is a man” “All men are mortal”). To the “doctrine of concepts” supporting the “doctrine of judgments” they then appended a “doctrine of syllogisms”, which classified inferences as good or bad, depending on the kinds of judgments they involved. (“Socrates is a man, and all men are mortal, *so* Socrates is mortal.”) This classical theory was a paradigm of the order of explanation that proceeds from the representational to the propositional dimensions of intentionality. In a radical break with tradition, Kant starts elsewhere. For him the fundamental intentional unity, the minimal unit of experience in the sense of sapient awareness is the *judgment* (proposition). For that is the minimal unit of *responsibility*. Concepts are to be understood top-down, by analyzing judgments (they are, he said “functions of judgment,” rules for judging), looking at what contribution they make to the responsibilities undertaken by those who bind themselves by those concepts in judgment (and intentional agency). He initiated an order of explanation that moves from the propositional to the representational dimensions of intentionality.

Pursuing that order of explanation in the context of his normative understanding of the propositional dimension of discursive intentionality led Kant to a normative account also of the representational dimension of discursive normativity. On the propositional side, the concept one has applied in judgment determines what one has made oneself responsible *for*. On the representational side, it determines what one has made oneself responsible *to*, in the sense of what sets the standard for assessments of the *correctness* of judgment. Kant sees that to treat something as a *representing*, as at least purporting to present something *represented*, is to acknowledge the *authority* of what is represented over assessments of the correctness of that representing. Discursive representation, too, is a normative phenomenon. And it is to be understood ultimately in terms of the contribution it makes to the normativity characteristic of propositional discursive intentionality.

Contemporary philosophical analyses of the normativity characteristic of discursive intentionality, along both propositional and representational dimensions, fall into two broad classes: social-practical and teleosemantic. Both are broadly *functionalist* approaches, in the

sense that they look to the role discursive intentional states play in some larger system in explaining the norms they are subject to. Teleosemantic theories derive norms (what ought to follow, how the representing ought to be) from selectionally, evolutionary, adaptive explanations of the advent of states and expressions that count as intentionally contentful (typically not just in the discursive, but also the practical sense) just in virtue of being governed by those norms. Ruth Millikan [b. 1933], for instance, defines Proper Function as that function that selectionally (counterfactually) explains the persistence of a feature or structure, in the sense that if such features had not in the past performed that function, it would not have persisted.⁵ Social practice theories date to Georg Hegel [1770-1831], who accepted Kant's insight into the normative character of discursive intentionality, but sought to naturalize the norms in question (which Kant had transcendentalized). He understood normative statuses, such as commitment, entitlement, responsibility, and authority, as instituted by practical normative attitudes. (Slogan: "All transcendental constitution is social institution."). On his account, genuine norms can only be instituted *socially*: as he put it, by "reciprocal recognition". The idea that discursive norms are to be understood as implicit in social practices was taken up from Hegel by the American pragmatists (C. S. Peirce [1839-1914], William James [1842-1910], and John Dewey [1859-1952]), and later on by Ludwig Wittgenstein [1889-1951], who had independently discovered the normative character of discursive content.

The idea is that social norms are instituted when practitioners take or treat performances *as* appropriate or inappropriate, take or treat each other *as* committed, entitled, responsible, authoritative, and so on. The pragmatist thought is that even if the norms in question are discursive norms, adopting the instituting normative attitudes might require only practical intentionality. Practically punishing or rewarding performances is one way of treating them as inappropriate or appropriate. So for instance hominids in a certain tribe might practically treat it as inappropriate for anyone to enter a certain hut without displaying a leaf from a rare tree, by beating with sticks anyone who attempts to do so. In virtue of the role they play in this practice, the leaves acquire the practical normative significance of hut-licenses. In more sophisticated cases, the reward or punishment might itself be an alteration in normative status, regardless of its

⁵ *Language, Thought, and other Biological Categories* [MIT Press, 1987].

actual reinforcing effect. So one might treat a performance as appropriate by giving the performer a hut-license leaf, even if he has no interest in entering the hut.

IV. An Inferential Approach to Discursive Propositional Intentional Content

What makes something a specifically *discursive* norm? Discursive norms are norms governing the application of concepts, paradigmatically in judgment. Discursive norms govern the deployment of judgeable, that is, propositional intentional contents. In the context of a commitment to pragmatism, this question becomes: what kind of knowing *how* (to do something) amounts to knowing (or believing) *that* (things are thus-and-so)? What is the decisive difference—the difference that *makes* the difference—between a parrot who can reliably differentially respond to the visible presence of red things, perhaps by uttering "Rawk! That's red," on the one hand, and a human observer who can respond to the same range of stimuli by claiming and judging *that* something is red? What is it that the sapient, discursively intentional observer knows how to do that the merely sentient, practically intentional parrot does not?

The important difference is, to be sure, a matter of a distinctive kind of *understanding* that the concept-user evinces. The pragmatist wants to know: what practical abilities does that understanding consist in? We have acknowledged already the *normative* difference: the observer's performance does, as the parrot's does not, express an endorsement, the acknowledgement of a commitment. The key additional point to understand is that the content endorsed, the content the sapient observer is committed to qualifies as a *conceptual* content (of which specifically *propositional* contents are a principal species) just insofar as it is situated in a space of other such contents to which it stands in relations of material *consequence* and *incompatibility*. The observer knows how to make inferences and so draw conclusions from his commitment: to determine what *else* he has committed himself to by the claim that the apple is red (for instance, that it is colored, that it is ripe...). He knows how to distinguish what is *evidence* for and against that claim, and what else that commitment rules out as incompatible (for instance, that it is not wholly green). The sapient practically understands his commitment as taking up a stance in a network of related possible commitments, which stand to one another in

rational relations of material consequence and incompatibility. He is making a move in a practice of giving and asking for *reasons*, in which one move has normative consequences for what others are obligatory, permitted, or prohibited.

Material inferential (and incompatibility) relations, by contrast to formal logical inferential and incompatibility relations, articulate the contents of non-logical concepts. These are inferences such as “A is to the West of B, so B is to the East of A,” “Lightning now, so thunder soon,” and “If the sample is copper, then it will conduct electricity.” Part of what one must do to count as understanding the contents of concepts such as East and West, lightning and thunder, copper and electrical conductor is to endorse inferences such as these. This is not to say that for each concept there is some meaning-constitutive set of material inferences one must endorse to count as understanding it. But if one makes *no* distinction, however partial and fallible, between material inferential and incompatibility relations that do and do not articulate the content of some concept, then one cannot count as a competent user of that concept.

Another way to get at the same point about the internal connection between *conceptual* contentfulness and *inferential* articulation is to consider the difference between *labeling* or *classifying* something and *describing* it. Any reliable differential responsive disposition imposes a classification on stimuli, distinguishing those that would from those that would not elicit a response of the given kind by the exercise of that reliable practical responsive capacity. The chunk of iron rusts in some environments and not others, the beam breaks under some loads and not others, the parrot squawks “Red!” in some situations and not others. What more is needed for such a performance to count not just as discriminating or labeling what elicits it, but also as describing it *as* red? The philosopher Wilfrid Sellars [1912-1989] offers the following inferentialist answer:

It is only because the expressions in terms of which we describe objects...locate these objects in a space of implications, that they describe at all, rather than merely label.⁶

If I discover that all the boxes in the attic I am charged with cleaning out have been labeled with red, yellow, or green stickers, all I learn is that those labeled with the same color share *some*

⁶ Pp. 306-307 (§107) in: Wilfrid Sellars: “Counterfactuals, Dispositions, and Causal Modalities” In *Minnesota Studies in the Philosophy of Science, Volume II: Concepts, Theories, and the Mind-Body Problem*, ed. Herbert Feigl, Michael Scriven, and Grover Maxwell (Minneapolis: University of Minnesota Press, 1958), p. 225-308.

property. To learn what they *mean* is to learn, for instance, that the owner put a red label on boxes to be discarded, green on those to be retained, and yellow on those that needed further sorting and decision. Once I know what *follows* from affixing one rather than another label, I can understand them not as *mere* labels, but as *descriptions* of the boxes to which they are applied. Description is classification with inferential *consequences*, either immediately practical (“to be discarded/examined/kept”) or for further classifications.

The inferentialist semantic claim is that what distinguishes specifically *discursive* (paradigmatically, but not exclusively, propositional) commitments is that their contents are articulated by the roles they play in material inferential and incompatibility relations. Grasping or understanding such contents is a kind of practical know-how: distinguishing in practice what follows from a given claimable or judgeable content, what it follows from, what would be evidence for it or against it, and what it would be evidence for or against. The practical inferential abilities to acknowledge the consequences of one’s commitments for further commitments (both those one is committed to and those one is precluded from) and to distinguish evidence that would and would not entitle one to those commitments are what distinguish sapient from mere sentient, creatures that exhibit discursive intentionality from those that exhibit only practical intentionality.

V. The Relation of Language and Thought in Discursive Intentionality

It is obvious that there can be practical intentionality without language. Can there be discursive intentionality in the absence of language? Modern philosophers from Descartes through Kant took it also to be obvious that propositionally contentful thoughts and beliefs both antedate and are intelligible apart from their linguistic expression, which they understood in terms of symbols whose meanings are inherited from those antecedent prelinguistic discursive states and episodes. More recently, H. P. Grice [1913-1988] extended this tradition, by understanding linguistic meaning in terms of speaker’s meaning, and speaker’s meaning in terms of the intention of a speaker to induce a belief in the audience by an utterance accompanied by the audience’s recognition that the utterance was produced with that very intention. Another

prominent line of thought in the area, due to Jerry Fodor [b. 1935], is the claim that public language is made possible by a language of thought, which is innate and so does not need to be learned.

A contrary order of explanation, identified with Wittgenstein among many others, gives explanatory priority to linguistic social practices in understanding discursive intentionality. Michael Dummett [b. 1925] forcefully expresses one of the consequences of this approach:

We have opposed throughout the view of assertion as the expression of an interior act of judgment; judgment, rather, is the interiorization of the external act of assertion.⁷

This way of turning the traditional explanatory strategy on its head is more extreme than is needed to acknowledge the crucial role of public language. Donald Davidson [1917-2003] claims that to be a believer in the discursive sense one must be an interpreter of the speech of others. But he also claims that:

Neither language nor thinking can be fully explained in terms of the other, and neither has conceptual priority. The two are, indeed, linked in the sense that each requires the other in order to be understood, but the linkage is not so complete that either suffices, even when reasonably reinforced, to explicate the other.⁸

Although Davidson shares some important motivations with Dummett's purely linguistic theory, in fact these two views illustrate an important difference between two ways in which one might give prominence to linguistic practice in thinking about discursive intentionality. Davidson's claim, by contrast to Dummett's, serves to epitomize a *relational* view of the significance of language for sapience: taking it that concept use is not intelligible in a context that does not include language use, but not insisting that linguistic practices can be made sense of without appeal at the same time to intentional states such as belief.

According to such relational views, the transition from mere sentience to sapience (from practical to discursive intentionality) is effected by coming into language: coming to participate in discursive, social, linguistic practices. The capacity to *think* in the discursive sense—that is,

⁷ Frege's *Philosophy of Language* [New York: Harper and Row, 1973], p. 362.

⁸ "Thought and Talk," in *Inquiries Into Truth and Interpretation* [New York: Oxford University Press, 1984], p. 156.

to have propositionally or conceptually contentful thoughts, to be able to think that things are thus-and-so (a matter of knowing *that*, not just knowing *how*)—and the capacity to *talk* arise and develop together. For Wittgenstein, the essentiality of public language to the capacity for individual thought is a consequence of the normativity of discursive intentionality. He endorsed a pragmatist order of explanation that understands discursive norms as in the first instance implicit in social practices (“uses, customs, institutions” as he put it).⁹ The capacity to make propositionally explicit claims and have conceptually contentful thoughts is intelligible only in the context of implicitly normative social linguistic practices.

VI. Putting Together a Social Normative Pragmatics and an Inferential Semantics for Discursive Intentionality

An inferentialist about discursive content who understands discursive norms as implicit in social linguistic practice and holds a relational view of the priority of language and thought will take it that the core of discursive intentionality is to be found in the role *declarative sentences* play in expressing *propositional* contents in speech acts of *assertion*. This connection between the *syntactic* category of declarative sentences, the *semantic* category of propositions, and the *pragmatic* category of assertions is the iron triangle of discursiveness. A pragmatist about the relations between them takes it that the syntactic and semantic elements are ultimately to be understood in terms of the pragmatic one. It is their role in the practice of assertion, of claiming that things are thus-and-so, that is appealed to in picking out declarative sentences and propositional contents. Propositional contents are what can both serve as and stand in need of reasons—that is, can perform the office both of premise and of conclusion in inferences. So the inferentialist pragmatist takes it that what distinguishes the speech act of assertion is its role in practices of giving and asking for *reasons*.

One way of putting together a social normative pragmatics and an inferential semantics for discursive intentionality is to think of linguistic practices in terms of *deontic scorekeeping*. Normative statuses show up as social statuses. The paradigmatic deontic status is *commitment*.

⁹ Ludwig Wittgenstein *Philosophical Investigations*, G.E.M Anscombe [Wiley-Blackwell, 3rd edition, 1991] §199.

The idea is that we should understand what one is doing in making an assertion is undertaking a distinctive kind of commitment: making a claim is staking a claim. If acquiring the status of being committed in the way standardly undertaken by assertively uttering the sentence *p* is to be significant, it must have consequences. The inferentialist says to look for *inferential* consequences (and antecedents): what *else* one becomes committed to by asserting *p* (what follows from *p*) and what would commit one to it (what it follows from). The pragmatist says to understand that in terms of what one is obliged (or permitted) to *do*, upon asserting *p*. To understand an assertional speech act is to know how to *keep score* on the commitments the speaker has undertaken by performing that act. In undertaking commitment to *p*, the asserter has *obliged* herself to acknowledge *other* commitments: those that follow from it. She has also *authorized* other interlocutors to attribute that commitment to her. Further, she has obliged herself to offer a *justification* (give reasons) for the claim, if her authority is suitably challenged. The idea is that exercising such inferentially articulated *authority* and fulfilling such inferentially articulated *responsibility* is what one must *do* (the task responsibilities one must carry out) in order to count as responsible for or committed—not now to do something, but to what in this social-practical scorekeeping context shows up as the propositional content *p*.

For such an idealized assertional practice to count as one of giving and asking for reasons, there must be a difference between commitments for which one *can* give a reason (so fulfilling one's justificatory task-responsibility) and those for which one *cannot*. That is, there must be a distinction between commitments to which an asserter is (rationally, inferentially, by one's evidence) *entitled*, and those to which the asserter is not entitled. So in practice to take or treat a performance *as* an assertion of a particular propositional content, other interlocutors must keep track not only of how that performance changes the score of what the asserter is committed to, but also what she (and others) are entitled to. Discursive scorekeeping requires attributing two sorts of deontic status: commitments and entitlements (to commitments), and knowing how different speech acts change the deontic "scores" of various interlocutors—who may become entitled to new commitments by relying on the authority of other asserters (to whom they can then defer their justificatory responsibility). This deontic scorekeeping story is a sketch of how

discursive intentionality is intelligible as emerging from exercises of practical intentionality that have the right normative and social structure.¹⁰

Scorekeepers acknowledging and attributing two kinds of normative deontic status, commitments and entitlements, can distinguish three kinds of practical consequential relations among them, which generate three flavors of inferential relations, and a relation of material incompatibility. Scorekeepers who take anyone who is entitled to p to be (prima facie) entitled to q thereby practically endorse a *permissive* inferential (probatively evidential) relation between p and q . This is a generalization, from the formal-logical to the contentful material case, of *inductive* inference. (The barometer is falling, so there will be a storm.) Scorekeepers who take anyone who is committed to p to be committed to q thereby practically endorse a *committive* inferential (dispositive evidential) relation between p and q . This is a generalization, from the formal-logical to the contentful material case, of *deductive* inference. (If the sample is pure copper, it will conduct electricity.) Scorekeepers who practically take or treat anyone who is committed to p *not* to be entitled to q , and vice versa, thereby treat the two claims they express as materially *incompatible*. (The plane figure cannot be both square and circular.) Scorekeepers for whom everything incompatible with q is incompatible with p thereby practically take or treat q as *incompatibility-entailed* by p . (Everything incompatible with Pedro being a mammal is incompatible with Pedro being a donkey, so his being a donkey in this sense entails his being a mammal.) These are modally robust, counterfactual-supporting entailments.

When an interlocutor makes an assertion by uttering p , scorekeepers take or treat him as also committed to committive consequences of p , withdraw attributed entitlements to any claims incompatible with p , and if they take it that he is also entitled to p , attribute further entitlements to its permissive consequences to him and to anyone in the audience not precluded by virtue of incompatible commitments. Adopting these practical deontic scorekeeping attitudes is what those who appreciate the practical significance of the speech act must *do* in order thereby to count as implicitly taking or treating the utterance as playing the functional role in virtue of which it expresses a *propositional* discursive content. Other uses of language are built on this

¹⁰ I develop this model further in *Making It Explicit: Reasoning, Representing, and Discursive Commitment* [Harvard University Press, 1994]—especially Chapter Three.

assertional-inferential core (the “downtown” of language), and make use of the conceptual contents conferred by it.

VI. Logic: the Organ of Semantic Self-Consciousness

According to this inferentialist social practical story about the structure of practical intentionality (knowing *how*, abilities) that adds up to discursive intentionality (knowing or believing *that* things are thus-and-so)—a story about pragmatics, or the *use* of language (the norms implicit in scorekeeping practices)—it is being practically taken or treated *as* standing in relations of material inference-and-incompatibility in virtue of which expressions come to have propositional discursive semantic content and so are able to make something explicit, in the sense of its being sayable, claimable, thinkable. Building on this kind of basic discursive (sapient) intentional practices and abilities, it is also possible for such practitioners to make propositionally explicit those normative material inferential and incompatibility relations, which are initially implicit in the practical attitudes discursive scorekeepers adopt to one another.

Most centrally, inferential (including *material* inferential) relations can be put in claimable (propositional, explicit) form by the use of *conditional* locutions. One can explicitly express one's endorsement of the inference from *p* to *q* by asserting "If *p* then *q*." Incompatibility relations can be made explicit using *negation* operators. One can explicitly express one's taking *p* to be incompatible with *q* by asserting "Not (*p*&*q*)."¹¹ Conditional and negation operators are *logical* vocabulary. (Indeed, versions of them suffice to define the classical propositional calculus.) The expressive role characteristic of logical vocabulary is to make explicit the material inferential and incompatibility relations in virtue of which non-logical vocabulary expresses the semantic content that it does. It is by playing the role they do in a network of such relations that expressions acquire the propositional content that makes possible the discursive, sapient awareness that consists in explicitly claiming or judging *that* things are

¹¹ I suppress here consideration of what modal operators (also logical vocabulary) make explicit. Incompatibility should really be rendered as “*Necessarily* not (*p* and *q*).” Incompatibility and modal operators are discussed in Chapter Five of *Between Saying and Doing* [op. cit.].

thus-and-so. Logical vocabulary makes possible explicit, discursive, sapient awareness of those very semantogenic material inferential and incompatibility relations. Logic is the organ of semantic self-consciousness.

On this account of the expressive role that demarcates vocabulary as distinctively logical, it is intelligible that there should be creatures that are *rational*, but not yet *logical*. To be rational is to engage in practices of giving and asking for reasons, that is, making inferentially articulated assertions and justifying them. To do that one must attribute and acknowledge commitments and entitlements, and practically keep track of their inferential relations along all three dimensions those two deontic statuses generate: permissive, committive, and incompatibility entailments. But one need not yet deploy specifically logical vocabulary, which permits one to make explicit and so be discursively aware of those material inferential and incompatibility relations. In being rational, one already knows how to do everything one needs to know how to do to introduce logical vocabulary. But until such semantically explicating vocabulary actually is deployed, rational creatures need not be semantically self-conscious, that is, logical creatures. We are not like that, but our hominid ancestors might have been.

VII. Pragmatic Social Normative Perspectives and the Representational Dimension of Discursive Semantic Content

Practically keeping track of inferentially-articulated commitments and entitlements (that is, engaging in discursive practices) requires distinguishing between the normative statuses one *attributes* (to another) and those one *acknowledges* (oneself). This distinction of social perspective between normative attitudes means that there are two points of view from which one can assess another's consequential commitments. For the auxiliary hypotheses or collateral premises one conjoins to another's avowed commitment to extract its consequences (whether permissive, committive, or incompatibility-entailed) can be drawn either from other commitments one *attributes* to that interlocutor, or from those one *undertakes* oneself. Suppose S attributes to A commitment to the claim "Benjamin Franklin was a printer," (perhaps on the basis of hearing A make that assertion). If S also attributes to A commitment to "Benjamin

Franklin is (=) the inventor of the lightning rod.” Then S should also attribute to A commitment to “The inventor of the lightning rod was a printer.” But suppose S, but not A, is committed to “Benjamin Franklin is (=) the inventor of bifocals.” Should S attribute to A commitment to “The inventor of bifocals was a printer”? Given the fact (as S takes it) that Franklin invented bifocals, that is indeed a consequence of A’s original claim. In the context of that fact, a claim about Benjamin Franklin is a claim about the inventor of bifocals, whether or not A realizes that. So in a genuine and important sense, A has, without knowing it, committed herself to the inventor of bifocals having been a printer. But that is a different sense from that in which A has committed herself to the inventor of the lightning rod having been a printer.

When the practical adoption of a normative attitude of attributing a commitment to another interlocutor is made propositionally explicit by the use of locutions that let one *say* what commitments one practically attributes to another, this difference in social perspective manifests itself in two different kinds of ascription of propositional attitude. Consequential commitments attributed solely on the basis of commitments the target would assert are ascribed *de dicto*. S can say “A claims (believes, is committed to the claim) *that* the inventor of the lightning rod was a printer.” Consequential commitments attributed partly on the basis of commitments the target would assert and partly by the use of collateral premises that the *attributor*, but *not* the target of the attribution, would assert are ascribed *de re*. S can say “A claims *of* the inventor of bifocals that he was a printer. In putting things this way, S marks that while he is *attributing* to A responsibility for the overall claim, S is himself *undertaking* responsibility for the substitution-inference licensed by the identity “Benjamin Franklin is the inventor of bifocals,” (commitment to which he does *not* attribute to A).¹²

Propositional attitude ascribing locutions, such as “claims” and “believes” let their users make explicit their practical normative scorekeeping attitudes of attributing commitments, that is, using such vocabulary empowers them to *say that* they adopt such attitudes, which otherwise remain implicit in what they practically do. Performing this expressive office with respect to social normative attitudes, on the side of pragmatics, marks them as another species of the same explicating genus as logical vocabulary, which does corresponding service on the semantic

¹² I discuss the distinction between propositional attitude ascriptions *de dicto* and *de re* in Chapter Eight of *Making It Explicit* [op. cit.], and Chapter Three of *Tales of the Mighty Dead: Historical Essays in the Metaphysics of Intentionality* [Harvard University Press, 2002].

side, by making explicit inferential commitments. What S is *doing* in making *de re* ascriptions is expressing the distinction of social perspective between commitments *attributed* (Ben Franklin was a printer) and those *undertaken* (Ben Franklin invented bifocals). But what one is *saying* is what the one to whom the commitments are ascribed was talking *about*. *De re* ascriptions of propositional attitude are the home language-game of *representational* locutions: the ones used to make explicit what one is talking or thinking *of* or *about*. What they make explicit is the representational dimension of discursive intentionality.

That representational dimension is always already implicit in the distinction of social perspective that is integral to keeping track of others' inferentially articulated commitments. For discursive deontic scorekeepers, players of the game of giving and asking for reasons, care about what follows from others' claims for two reasons. They care about the consequential commitments that would be ascribed *de dicto* because they want to know what else the target would endorse, and what she will do based on the commitments she acknowledges. They care about the consequential commitments that would be ascribed *de re* because they want to extract *information* from the claims of others—that is, premises that the attributor can use in his own inferences. If S attributes to A the intention to shoot a deer and the belief that the tawny creature in front of her is a deer, the *de dicto* ascription “A believes that the tawny creature in front of her is a deer, the shooting of which would fulfill her intention,” S will predict that A will shoot. If S, but not A, believes that the tawny creature in front of A is (=) a cow, then S's *de re* ascription “A believes *of* the cow in front of her *that* it is a deer, the shooting of which would fulfill her intention,” S will predict that the result of A's action will be the shooting of a cow. That is an inference that S is in a position to extract from A's avowed commitments, even though that information is not available to A. Keeping track of what premises are available for the reasoning of others and what premises are available for our own reasoning is what we are doing when we talk or think about what we are talking or thinking *about*: the representational dimension of discursive intentionality.

END

[6987 words, excluding footnotes]