



Photo © Mick Canning

Panoramas of Mind and Meaning

Robert Brandom

Contents:

1. Kantian Lessons about Mind, Meaning, and Rationality	5
2. From German Idealism to American Pragmatism, and Back	38
3. Reason, Expression, and the Philosophic Enterprise	80
4. Intentionality and Language: A Normative, Pragmatist, Inferentialist Approach	105
5. Semantic Inferentialism and Logical Expressivism	130
6. How Analytic Philosophy Has Failed Cognitive Science	168
7. Reason, Genealogy, and the Hermeneutics of Magnanimity	203

Preface:

This is a collection of previously published essays written for more general audiences than I usually address: for interested academics from all fields, for nonphilosophers, and for philosophers who might not share my somewhat specialized interests. The aim in each case is to introduce serious, substantive, controversial philosophical ideas without presupposing a lot of background or compromising either the integrity of those ideas or the precision of expression they deserve and require. Taken together, these essays present a more or less unified perspective on some of the largest issues concerning us essentially rational, discursive, normative beings *as* essentially rational, discursive, normative beings.

Previous Publications:

1. "Kantian Lessons About Mind, Meaning, and Rationality":
Southern Journal of Philosophy, Vol 44 Supplement, 2006, pp. 49-71.
Also in *Philosophical Topics*, Vol 34, Nos. 1 & 2, 2006, pp. 1-20.
2. "From German Idealism to American Pragmatism, and Back":
In Robert Brandom, *Perspectives on Pragmatism: Classical, Recent, and Contemporary* [Harvard University Press, 2011], pp. 1-34.
3. "Reason, Expression, and the Philosophic Enterprise" in *What Is Philosophy?*, C.P. Ragland and Sarah Heidt (eds.) [Yale University Press, 2001], pp. 74-95. Reprinted in Robert Brandom, *Reason in Philosophy: Animating Ideas* [Harvard University Press, 2009], pp. 111-129.
4. "Intentionality and Language: A Normative, Pragmatist, Inferentialist Approach" in N.J. Enfield and Paul Kockelman (ed.s) *The Cambridge Handbook of Linguistic Anthropology* [Cambridge University Press, 2014], Chapter 14.
5. "Semantic Inferentialism and Logical Expressivism" in Robert Brandom *Articulating Reasons* [Harvard University Press, 2001] pp. 45-78.
6. "How Analytic Philosophy Has Failed Cognitive Science" *Critique and Humanism*, Vol. 31, No. 1, 2010, pp. 151-174. Reprinted in Robert Brandom, *Reason in Philosophy: Animating Ideas* [Harvard University Press, 2009], pp. 197-225.
7. "Reason, Genealogy, and the Hermeneutics of Magnanimity" the 2013 Howison lecture at the University of California/Berkeley. Video available online at <http://www.uctv.tv/shows/Reason-Genealogy-and-the-Hermeneutics-of-Magnanimity-with-Robert-Brandom-25074> and published in German as "Den Abgrund reflektieren" in *Westend: Neue Zeitschrift für Sozialforschung* **12** (2015), 1.

Panoramas of Mind and Meaning

Chapter One:

Kantian Lessons about Mind, Meaning, and Rationality¹

Kant revolutionized our thinking about what it is to have a mind. Some of what seem to me to be among the most important lessons he taught us are often not yet sufficiently understood, however. I think this is partly because they are often not themes that Kant himself explicitly emphasized. To appreciate these ideas, one must look primarily at what he *does*, rather than at what he *says* about what he is doing. For instance, one revolutionary conceptual transformation Kant focuses on is his “Copernican Revolution”: assignment of responsibility for some structural features of knowledge to the nature of the activities of knowing subjects rather than to the nature of the objects known. While this is, of course, an important aspect of his view, as I understand things it is a relatively late-coming move; it occurs significantly downstream from his most radical and important innovations, whose significance owes nothing to this subsequent, optional way of developing them. Indeed, I will say nothing at all about Kant’s transcendental idealism (important as that view was to his way of setting out his achievement). For I want to focus on revolutionary moves that happen off-stage, largely before the first Critique even begins, but which seem to me to form a crucial backdrop and stage for that performance. So I will sketch

¹ The material in this chapter was substantially expanded in my 2007 Woodbridge Lectures at Columbia University, published as *Animating Ideas of Idealism: A Semantic Sonata in Kant and Hegel*, Part One of *Reason in Philosophy* [Harvard University Press, 2009].

here in very broad terms some kantian lessons that it seems to me most important for us to keep in mind in our own thinking about mind, meaning, and rationality.

1. *From Epistemology to Semantics*

Descartes gives philosophical thought about the mind an *epistemological* turn by using the character of our knowledge of them to distinguish minds from bodies. Thoughts are understood as distinguished by their epistemic transparency and incorrigibility: their immunity respectively to ignorance and error. The rest of the world is that about which we can be ignorant or in error. Our thoughts can *misrepresent* their objects, or represent them incompletely. That is why our knowledge of the merely represented world is fallible and incomplete. But on pain of an infinite regress, those thoughts must themselves be understood as known *immediately*, by being *had* rather than by being *represented* in their turn—hence their privileged epistemic status. At the most basic level, however, Descartes just takes it for granted that the world comes in two flavors: stuff that by nature represents and stuff that is by nature represented. The representing stuff is *intrinsically* “*tanquam rem*”, as if of things. The question he focuses on is what reason we have to think that things are in fact as we represent them to be. The fundamental Cartesian problematic is accordingly to explain the possibility of *knowledge*, that is of beliefs about how things are outside the mind that are both *true* and *justified*. How can we show that things *really are* as they *appear* to us to be—that is, how can we *justify* that claim?

By contrast, Kant gives philosophical thought about the mind a *semantic* turn by shifting the center of attention from truth and justification to the nature of representation itself. He replaces concern with justifying claims to representational *success* by concern with understanding representational *purport*. His problem is not in the first instance showing that reality at least often is as it appears, but understanding what it is for things so much as to *appear* to be one way rather than another. Kant wants to know what it is for mental states to *be*, or to *appear to us* to be, to function for us as, *representings* of represented objects. Kant sees that the *epistemological* question has *semantic* presuppositions. The issue for him is not *knowledge*, but *intentionality*.

These projects can each be thought of as responding to the threat of *skepticism*. But the kinds of skepticism addressed are quite different. Descartes worries about responding to the threat of *epistemological* skepticism: things may not in fact be at all as we take them to be. Or at least, we can't *show* that they are. Kant worries about responding to the threat of a deeper and more radical *semantic* skepticism. This is the claim that the very idea of our mental states *purporting* to specify how things are is unintelligible. Kant's most basic transcendental question does not, as his own characterization of his project suggests, concern the condition of the possibility of synthetic knowledge *a priori*, but the conditions of the intelligibility of representational *objectivity*: of states or episodes that answer for their correctness to how it is with the objects they represent.

In asking this question, Kant moves to an issue that is clearly conceptually prior to the one that is central for Descartes. And this move is not of merely historical interest. The principal argument of Sellars's masterwork *Empiricism and the Philosophy of Mind* is that the soft underbelly of both traditional and logical empiricism is their implicit semantics. Broadly Cartesian foundationalism depends on there being a *semantically autonomous* stratum of thought—what is 'given', *both* semantically *and* epistemologically. It is this *semantic* givenness that Sellars ultimately takes issue with. So Sellars offers Kantian *semantic* arguments against the *epistemological* Myth of the Given. More specifically, Sellars argues that there cannot be an autonomous language game—one that can be played though no other is—that consists entirely of making non-inferential reports. Unless some claims (endorsements) can be made as the conclusions of *inferences*, none of them can count as *conceptually* contentful, in the sense required for them even potentially to offer *evidence* or *justification* for further conclusions. That is, nothing that cannot serve as the *conclusion* of inference can serve as the *premise* for one. We'll see further along that this, too, is a Kantian theme. But for now I'm not concerned to say why one might think Sellars is *right* on this point (though I think he is)—only that he is developing Kant's response to Descartes.

2. A Semantic Response to Epistemological Skepticism

So far I've said that Kant points out that before one worries about whether and under what circumstances mental representation is and can be known to be generically *successful* one needs to think hard about what it is for something to be taken or treated as, to have the practical or

functional significance of, a *representation* at all. In classical epistemological terms, one must understand what it is to *believe that* things are thus-and-so in order then to try to understand what it is for such a belief to be *true* or *justified*. But Kant in fact commits himself to something much stronger than this. He thinks that any adequate answer to the *semantic* skeptic will in fact be an adequate answer also to the *epistemological* skeptic. Specifying the conditions under which there can be representings at all will settle it that *some*, indeed, *many* of them must be *true*. This is the idea behind his “Refutation of Idealism”: once one has seen what is presupposed by representational *purport*, one will see that it includes a substantial degree of representational *success*. Unless we are to a large extent *right* about how things are, we can’t make sense even of our being *wrong* in special cases about how things are.

This is an exceptionally bold claim. Once again, it is echoed and developed in our own time. Davidson argues that his interpretivist semantic methodology underwrites a principle of charity, which in turn has the consequence that in order properly to understand creatures as meaning or believing anything at all—as having a mind, or being rational—we must take it that most of their beliefs are *true*. Again, the lesson that Putnam draws from his analysis of the ‘thought’s of brains-in-vats is that *semantic* externalism has the *epistemological* consequence of ruling out radical skepticism: for the brains-in-vats to have thoughts *representing* an external world at all, they must be sufficiently in contact with it to have many *true* beliefs about it.

Looking back from the vantage point these arguments have secured for us, it seems to me that we can see that semantic arguments for this sort of epistemological conclusion must proceed in two

stages. First, one must argue that for any conceptual contents of the sort that represent things as being one way or another to be *entertained*, some others must be *endorsed*. Considering how some things *might* be (as it were, merely *representing* them as being thus-and-so) is not intelligible in total independence of taking other things *actually* to be one way or another. Such a Quinean thought contradicts methodologies (for instance, those of Dretske and Fodor) that depend on a layer-cake picture, according to which *first* one tells an autonomous semantic story about the possession of representational content by certain episodes, and only *then* adds on a story about what it is for some of those representations to play the functional role of being *accepted*, that is, to be taken to be *correct* or *successful* representations: to be *endorsed*, and not merely *entertained*. Frege notoriously treats merely entertaining a proposition as a speech-or-thought act wholly derivative from and dependent upon actually endorsing one (taking it to be true) in judgment or assertion. For him, merely entertaining a proposition is just endorsing various conditionals in which it appears as antecedent or consequent—and thereby exploring the circumstances under which it would be true, and the consequences that would ensue were it true.

It is not my purpose here to argue for one or the other of these ways of construing things. It is my purpose to point out that where the tradition Descartes inaugurated took it for granted that one could make autonomous sense of a mind as merely *entertaining* various fully contentful ideas or representations, and only then consider the “act of the will” that is plumping for or endorsing some of them (what in the contemporary context shows up as “putting representations in the belief box”), Kant’s idea of the understanding as in the first instance a function of *judgment*—concepts as intelligible only in terms of their contribution to the activity of judging—offers a radically different approach.

The second step in a semantic argument against epistemological skepticism would then have to be a justification of the claim that we cannot make sense of a whole constellation of representations, some of which are merely entertained and others of which are endorsed, unless we take it that many, perhaps most of the representations endorsed are correct or successful. Thus Davidson claims that local error is intelligibly attributable only against the background of an attribution of a good global grip on how things really are. And Putnam claims in effect that securing reference to natural kinds and individuals requires many true collateral beliefs about them.

An argument along these lines may or may not work. But Kant's idea that one could show on *semantic* grounds that we have sufficient reason to reject global *epistemological* skepticism about the truth and justification of our beliefs in general is a deep and radically original one. On top of Kant's semantic *transformation* of philosophic problematics, from epistemological to semantic, he builds a semantic explanatory *aspiration*: that resolving the semantic problematic would resolve the epistemological one. At the end of this essay, I'll say something about how this general aspiration is worked out and applied to the epistemological predicament Kant saw Hume as leaving us in with respect to modal and normative concepts.

3. *Force and Content*

In making the point about how one might take responses to *semantic* skepticism to bear on *epistemological* skepticism, I invoked a distinction between two sorts of things one could *do* with representations: merely *entertain* them, or further, *endorse* them. But I don't think that in the end this is the most helpful way to think about the thought that underlies Kant's views in this vicinity. The thought that motivates Kant here is the importance of distinguishing between what Frege calls 'force' and 'content'. This is the distinction between what one is *doing* in endorsing a claim—taking it to be true, whether internally by judging or externally by asserting—on the one hand, and what one thereby endorses, on the other. That is, Kant's practice depends on distinguishing between the two sides of what Sellars called "the notorious 'ing'/'ed' ambiguity": between judgment as the act of *judging* and as the content *judged*. (A distinction of cardinal importance, for instance, in sorting out Berkeley's confusions, conflations, and equivocations regarding 'experience' in the sense of *experiencings* and what is *experienced*.) The Kant-Frege claim is that to think of merely entertaining a representation as something one can *do* is to fail to appreciate the distinction between judging and what is judged (between force and content).

The tradition Kant inherited understood judging as *predicating*: classifying something particular as being of some general kind, applying a universal concept to a particular one. Although Kant continues to use the traditional language (thereby distracting attention from the radical break he is making from that tradition on this point), he sees that this will not do. His table of the forms of judgments includes *conditional*, *disjunctive*, *negative* and *modal* judgements, none of which kinds is happily assimilated to the predicational-classificatory model. The underlying thought is not made fully explicit until Frege. In the traditional theory, the notion of predication is being asked to do two incompatible jobs. On the one hand, it serves as a structural way of building up

new judgeable contents. On the other hand, it is thought of as a kind of doing that has the significance of endorsing such contents. The collision between these two senses in which predication is an ‘operation’ is clearest when one thinks about judgeable contents appearing as unasserted (unendorsed) components of more complex sentences (judgments). The conditional is a paradigm. When I assert “If *Pa* then *Pb*,” I have *not* asserted *Pa*. Have I predicated *P* of *a*? If so, then predication does not amount to endorsement: predicating is not judging. If not, then it looks as though there is an equivocation when I detach from the conditional, reasoning:

If *Pa* then *Pb*

Pa

So: *Pb*

For the second premise *is* a predication, and the antecedent of the first premise is *not* a predication.

Geach picks up this Kant-Frege point, using it in his masterful gem-like essay “Ascriptivism” to argue against emotivist semantic analyses of terms of moral evaluation. His target is theories that understand the normative significance of terms such as ‘good’ not as part of the *content* of what is said about an act, not as specifying a characteristic that is being attributed, but rather as marking the *force* of the speech act. Calling something good is thought of as *doing* something distinctive: commending. Geach first asks what the limits of this ploy are. He points to the archaic English verb “to macarize”, meaning to characterize someone as happy. Does the possibility of understanding calling someone happy as macarizing her mean that happiness is not a property being invoked in specifying the content of the claim that someone is happy, because in saying that we are really *doing* something else, performing the special speech act of macarizing? If we can do that with ‘happy’, why not with ‘mass’ or ‘red’? What are the rules of

this game? He then suggests the embedding test: look to see if an expression can be used to construct a judgeable content that is *not* directly used to perform a speech act, paradigmatically in the antecedent of a conditional. Because imperatival force *is* grammatically marked, we cannot say:

*“If shut the door, then....”

But we *can* say things like “If he is happy, then I am glad,” and “If that is a good thing to do, then you have reason to do it.” In the first of these, I have *not* macarized anyone, and in the second, I have not commended any action. So the terms ‘good’ and ‘happy’ contribute to the specification of content, and are not to be understood as mere force indicators.

Worrying about compound forms of judgment containing unendorsed judgeable contents as components required Kant to distinguish the operations by which such contents are constructed from the activity of endorsing the results of those operations.

4. Normativity and Force

For this reason, Kant could not take over the traditional classificatory theory of consciousness, which depends on understanding judging as predicating. But what can go in its place? Here is perhaps Kant’s deepest and most original idea, the axis around which I see all of his thought as revolving. What distinguishes judging and intentional doing from the activities of non-sapient creatures is not that they involve some special sort of mental processes, but that they are things knowers and agents are in a distinctive way *responsible* for. Judging and acting involve *commitments*. They are *endorsements*, exercises of *authority*. Responsibility, commitment,

endorsement, authority—these are all *normative* notions. Judgments and actions make knowers and agents liable to characteristic kinds of *normative* assessment. Kant's most basic idea is that minded creatures are to be distinguished from un-minded ones not by a matter-of-fact ontological distinction (the presence of mind-stuff), but by a normative *deontological* one. This is his *normative characterization* of the mental.

Drawing on a jurisprudential tradition that includes Grotius, Pufendorf, and Crusius, Kant talks about norms in the form of *rules*. Judging and acting—endorsing claims and maxims, committing ourselves as to what is or shall be true—is binding ourselves by norms. It is making ourselves subject to assessment according to rules that articulate the *contents* of those commitments. Those norms, those rules, he calls 'concepts'. In a strict sense, all a Kantian subject can do *is* apply concepts, either theoretically, in judging, or practically, in acting. Discursive, that is to say, concept-mongering creatures, are normative creatures—creatures who live, and move, and have their being in a normative space.

It follows that the most urgent philosophical task is to understand the nature of this normativity, the bindingness or validity (*Verbindlichkeit*, *Gültigkeit*) of conceptual norms. For Descartes, the question was how to think about our grip on our concepts, thoughts, or ideas (Is it clear? Is it distinct?). For Kant the question is rather how to understand their grip on us: the conditions of the intelligibility of our being bound by conceptual norms.

5. *Pragmatism about the Relations between Force and Content*

This master idea has some of Kant's most characteristic innovations as relatively immediate consequences. The logical tradition that understood judging as predicating did so as part of an order of semantic explanation that starts with concepts or terms, particular and general, advances on that basis to an understanding of judgements (judgeables) as applications of general to particular terms, and builds on that basis an account of inferences or consequences, construed syllogistically in terms of the sort of predication or classification exhibited by the judgments that appear as premises and conclusions. In a radical break with this tradition, Kant takes the whole judgment to be the conceptually and explanatorily basic unit at once of meaning, cognition, awareness, and experience. Concepts and their contents are to be understood only in terms of the contribution they make to judgments: concepts are functions of judgment. Kant adopts this semantic order of explanation because judgments are the minimal units of *responsibility*—the smallest semantic items that can express *commitments*. The semantic primacy of the propositional is a consequence of the central role he accords to the *normative* significance of our conceptually articulated doings. In Frege this thought shows up as the claim that judgeable contents are the smallest units to which pragmatic force can attach: paradigmatically, assertional force. In the later Wittgenstein, it shows up as the claim that sentences are the smallest linguistic units with which one can make a move in the language game.

In conditioning the semantic account of content on the pragmatic account of force—the way the story about what is endorsed is shaped by the story about what endorsing is—Kant exhibits a kind of methodological *pragmatism*. In this sense, that pragmatism consists not in the explanatory privileging of *practical* discursive activity over *theoretical* discursive activity, but in the explanatory privileging of *force* over *content*, within *both* the theoretical and the practical domains. Kant’s idea is that his *normative characterization* of mental *activity*—understanding judging and acting as *endorsing*, taking responsibility for, committing oneself to, some content—is the place to start in understanding and explaining the nature of the representational, object-presenting judgeable *contents* of those judgments. This explanatory strategy is Kant’s *pragmatic turn*.

It is this order of explanation that is responsible for the most general features of Kant’s account of the form of judgment. The subjective form of judgment is the ‘I think’ that can accompany all our judgments, and so, in its pure formality, is the emptiest of all representations. Thought of in terms of the normative pragmatics of judgment, it is the mark of *who* is responsible for the judgment. (A corresponding point applies to the endorsement of practical maxims.) The transcendental unity of apperception is ‘transcendental’ because the sorting of endorsements into co-responsibility classes is a basic condition of the normative significance of commitments. Committing myself to the animal being a fox, or to driving you to the airport tomorrow morning normatively preclude *me* from committing myself to its being a rabbit, or to *my* sleeping in tomorrow, but they do not in the same way constrain the commitments others might undertake.

The objective form of judgment is “the object=X” to which judgments always, by their very form as judgments, make implicit reference. Thought of in terms of the normative pragmatics of judgment, it is the mark of what one has made oneself responsible *to* by making a judgment. It expresses the objectivity of judgments, in the sense of their having intentional objects: what they purport to represent. The understanding of the intentional directedness of judgments—the fact that they are *about* something—is through-and-through a *normative* one. What the judgment is about is the object that determines the *correctness* of the commitment one has undertaken by endorsing it. (On the practical side, it is normative assessments of the *success* of an action for which the object to which one has made oneself responsible by endorsing a maxim must be addressed.) In endorsing a judgment one has made oneself liable to distinctive kinds of normative assessment. What one is thinking and talking *about* is what plays a special role, exercises a special sort of *authority* in such assessments. Representing something, talking *about* or thinking *of* it, is acknowledging its semantic *authority* over the correctness of the commitments one is making in judging. Representational purport is a normative phenomenon. Representational content is to be understood in terms of it.

6. *Reasons and Content*

Intentionality—semantic contentfulness—comes in two flavors: ‘of’-intentionality and ‘that’-intentionality. The first, or *representational* dimension, is semantic directedness at objects: what one is thinking *of* or talking *about*. The second, or *expressive* dimension, concerns the *content* of

our thought and talk: what one is thinking or saying (about what one is thinking or talking about). So one can think *of* or *about* foxes, *that* they are nocturnal omnivores. What falls within the scope of the ‘of’ in such a specification is a term, while what follows the ‘that’ in such phrases as “I think (or John thinks) *that* foxes are nocturnal omnivores,” is a declarative sentence. The pre-Kantian early modern philosophical tradition took it for granted that one ought first to offer an independent account of representational, ‘of’-intentionality, of what it is to represent something, and only then, on that basis to explain expressive, ‘that’-intentionality, what it is to judge or claim *that* things are thus-and-so. It is part and parcel of Kant’s semantic revolution to reverse that order of explanation.

That is to say that just as he needs a new and different idea about what one is doing in judging, on the pragmatic side of force, so he needs a new and different idea about what that force attaches to or is invested in, on the semantic side of content. His thought that judging is taking responsibility, committing oneself, requires a corresponding characterization of what one thereby becomes responsible *for*, commits oneself *to*. The contents of judgments are articulated by *concepts*. The conceptual faculty, the understanding, is the faculty of judgment. Concepts articulate the contents of judgments by determining what one would make oneself responsible for, what one would be committing oneself to, were one to endorse those contents.

I think that at this point Kant wheels in a Leibnizian idea: concepts are in the first instance rules that express what is a *reason* for what. The concepts being applied determine what *follows* from a given claim(able), hence what (else) one would have committed oneself to or made oneself

responsible for by endorsing it. They determine what counts as rational *evidence for* or *justification of* a claim(able) content, hence would count as a *reason for* endorsing it. An essential element of what one is responsible for in endorsing a claim or a maxim is having *reasons* for doing so. That is part of the responsibility that goes with investing one's authority in the claim or maxim. Norms must have content, and the concepts that articulate those contents are rules specifying what is a reason for what. As *normative* creatures, we are *rational* creatures—not in the sense that we always or even generally do what we ought, or that we usually have good reasons for doing what we do, but in the sense that whether we do or not, we are always liable to normative *assessment* concerning our reasons for doing what we do, or thinking as we do. However *sensitive* we are in fact on any particular occasion to the normative force of reasons (that peculiar force, at once compulsory and yet not always compelling, that so fascinated and puzzled the ancient Greek philosophers), we are the kind of creatures we are—knowers and agents, creatures whose world is structured by the *commitments* we undertake—only because we are always liable to normative assessments of our *reasons*. *Discursive* creatures are those bound by *conceptual*, that is to say, *inferentially* articulated norms. It is at this level that Kant applies the lessons he learned from his *rationalist* predecessors.

To complete the semantic story, at this point an account is needed of the relation between the two kinds of intentionality: representational and expressive. That is, Kant must explain the contribution that the objects we become responsible *to* in judgment make to what we thereby become responsible *for*. This is the task to which he devotes the bulk of his efforts in the *Transcendental Analytic* and the *Transcendental Aesthetic*. It is in the service of this project that he introduces the faculty of sensuous intuition, the faculty by which particular objects are

understood to be empirically given to us. His story about how to understand sensuous receptivity in terms of its role in or contribution to the contents of the concepts applied in empirical judgment (and hence, experience) is intricate, instructive, and fascinating. I'm not going to say anything about it here, for two reasons. First, a fundamental structural element of his story depends on lining up as essential dimensions of the intuition/concept distinction what I take to be three quite different (indeed, orthogonal) distinctions—that between receptivity and spontaneity, that between singular terms and predicates, and that between unrepeatable and repeatable representations (tokens and types). (It is startling to see the thinker who marked so carefully the distinction between representations of relations and relations of representations run together representations of particularity and particularity of representations.) In order to extract the important insights that are in play in his discussion, I think one must divide through by this mistaken assimilation. And that is no easy or straightforward task. More generally, however, my concern in this essay is to emphasize the radical and revolutionary conceptual shifts that Kant makes as part of the stage-setting for his assault on the problem he puts in the foreground of his text—elements that are in danger of remaining unnoticed in the background, but which may in fact constitute his best claim to contemporary philosophical attention and admiration.

Be that as it may, it is at this subsidiary explanatory level that I see Kant applying—for better or for worse²—the lessons he learned from his *empiricist* predecessors. The semantic explanatory strategy of understanding and explaining representational 'of'-intentional content in terms of expressive 'that'-intentional content is Kant's *propositional* turn in semantics. I see it as a

² John McDowell insists on the former, while I am inclined to the latter assessment.

consequence of his normative and pragmatic turns. In terms of later developments, we can see it as a question of the relative explanatory priority of the notions of the sense *expressed* by a claim and the object *represented* by a singular term. With the wisdom of hindsight vouchsafed us by Frege's analysis (still opaque to Russell), we can see that two issues still remain to be disentangled here: the distinction between the content associated with declarative sentences and that associated with singular terms, and the distinction between sense and reference. Still, in this area Kant has once again not only made a crucial distinction, but on principled grounds endorsed a bold, unprecedented, and promising order of explanation.

7. *Freedom*

Against the background of this set of ideas about normativity and rationality—which is to say his new ways of understanding pragmatic force and semantic content, and their relations to one another—Kant introduces a radically novel conception of *freedom*. Before Kant, freedom had traditionally, though not exclusively, been understood in *negative* terms: as freedom *from* some kind of constraint. He revolutionized our thought by introducing a special idea of *positive* freedom: freedom *to* do something. Positive freedom is generically a kind of ability or practical capacity. Even if I am not tied up, threatened, or otherwise restrained from playing the Minute Waltz, and hence am in the negative sense entirely free to do so, I am nonetheless not free to play it in the positive sense if I don't have a piano available, or do not know how to play the one that is available.

Kant's specific conception of positive freedom is normative. Being free is being able to adopt normative statuses, paradigmatically, to commit oneself, to undertake responsibilities. It is the capacity to bind oneself by conceptual norms, in judgment and action. This is exercising a certain kind of inferentially articulated authority—a kind that comes with a correlative rational responsibility to have reasons for one's endorsements. To use an example suggested by Kant's metaphor in "What is Enlightenment?" consider what happens when young people achieve their legal majority. Suddenly they can enter into contracts, and so legally bind themselves. Hence they can do things such as borrow money, start businesses, and take out mortgages. This change of normative status involves a huge increase in positive freedom. The difference between discursive creatures and non-discursive ones is likewise to be understood in terms of the sort of normative positive *freedom* exhibited by the concept-users. On this account, being free is not only *compatible* with constraint by norms, it *consists* in constraint by norms. Since the norms are *conceptual* norms, their content is articulated by *reasons*. Positive normative freedom is the capacity to act for reasons, not in a causal sense, but in the normative sense of the ability to bind oneself by norms that make one liable to assessment as to one's reasons.

This constellation of ideas about normativity, reason, and freedom is, I think, what Heidegger means when he talks about "the dignity and spiritual greatness of German Idealism."

8. *Autonomy and Normativity*

One of the permanent intellectual achievements and great philosophical legacies of the Enlightenment—and perhaps the greatest contribution philosophers have made to the wider culture—was the development of secular conceptions of legal, political, and moral normativity. In the place of traditional appeals to authority derived ultimately from divine commands, thought of as ontologically based upon the status of the heavenly lord as creator of those he commands, Enlightenment philosophers conceived of kinds of responsibility and authority (commitment and entitlement) that derived from the practical attitudes of human beings. So for instance in social contract theories of political obligation, normative statuses are thought of as instituted by the intent of individuals to bind themselves, on the model of promising or entering into a contract. Political authority is understood as ultimately derived from its (perhaps only implicit) *acknowledgment* by those over whom it is exercised. Following Rousseau, Kant radicalizes this line of thought, developing on its basis a new criterion of demarcation for the normative, in terms of *autonomy*. This is the idea that we are genuinely normatively constrained only by rules we constrain *ourselves* by, that we adopt and acknowledge *as* binding on us. The difference between non-normative *compulsion* and normative *authority* is that we are genuinely normatively responsible only to what we acknowledge as authoritative. In the end, we can only bind ourselves, in the sense that we are only bound by the results of exercises of our freedom: self-bindings, commitments we have undertaken by acknowledging them.

The acknowledgement of authority may be merely implicit, as when Kant argues that in acknowledging others as concept users we are implicitly also acknowledging a commitment not to treat their concept-using activities as mere means to our own ends. That is, there can be background commitments that are part of the implicit structure of rationality and normativity as

such. But even in these cases, the source of our normative *statuses* is understood to lie in our normative *attitudes*. Merely natural creatures are bound only by rules in the form of laws whose bindingness is not at all conditioned by their acknowledgment of those rules as binding on them. Normatively free, rational creatures are also bound by norms, which is to say by rules that are binding only insofar as they are acknowledged as binding by those creatures. As Kant says, we are bound not just by rules, but by *conceptions* of rules.

9. *The Force and Content of Conceptual Norms*

It is important to notice that this picture requires the strict conceptual separation of the *content* of norms from their normative *force*. The Kant-Rousseau autonomy understanding of the nature of the force or bindingness of norms is that it is always self-binding. Only we ourselves can *normatively* bind ourselves. It is in the end up to us what we are committed to and responsible for (though acknowledging *any* conceptual commitments may involve further implicit rationality- and intentionality-structural commitments). If not only the normative *force*, but also the *contents* of those commitments were *also* up to us, then, to paraphrase Wittgenstein, “whatever seems right to us would be right” and talk of what is right or wrong could get no intelligible grip: no norm would have been brought to bear, no genuine commitment undertaken. Put another way, autonomy, binding oneself by a norm, rule, or law, has two components, corresponding to ‘autos’ and ‘nomos’. One must bind *oneself*, but one must also *bind* oneself. If not only *that* one is bound by a certain norm, but also *what* that norm involves—what is correct or incorrect according to it—is up to the one endorsing it, the notion that one is *bound*, that a

distinction has been put in place between what is correct and incorrect according to that norm goes missing. The attitude-dependence of normative *force*, which is what the autonomy thesis asserts, is intelligible only in a context in which the boundaries of the *content*—what I acknowledge as constraining me and by that acknowledgment *make* into a normative *constraint* on me in the sense of opening myself up to normative *assessments* according to it—are *not* in the same way attitude dependent. That is a condition of making the notion of normative constraint intelligible.

Kant secures this necessary division of labor by appeal to *concepts*, as rules that determine what is a reason for what, and *so* what falls under the concepts so articulated. (If being malleable is a conclusive consequence of being gold, then only malleable particulars can fall under the concept gold.) His picture of empirical activity as consisting in the application of concepts—of judging and acting as consisting in the endorsement of propositions and maxims—strictly separates the contents endorsed from the acts of endorsing them. The latter is our responsibility, the former is not. This does not require that the constitution of conceptual contents be wholly independent of our activity. Kant in fact sees “judgments of reflection” as playing a crucial role in it. It requires only that each empirical (“determinate”) judgment be made in a context in which already determinately contentful concepts are available as candidates for application. The judging or acting empirical consciousness always already has available a stable of completely determinate concepts. Its function is to choose among them, picking which ones to invest its authority in by applying to objects, hence which conceptually articulated responsibility to assume, which discursive commitments to undertake. Judging that what I see ahead is a *dog*—applying that concept in perceptual judgment—may initially be successfully integratable into my

transcendental unity of apperception, in that it is not incompatible with any of my other commitments. But subsequent empirical experience may normatively require me to withdraw that characterization, and apply instead the concept *fox*. That is my activity and my responsibility. But what other judgments are compatible with somethings being a dog or a fox is *not* at that point up to me. It is settled by the contents of those concepts, by the particular rules I can choose to apply.

In taking this line, Kant is adopting a characteristic rationalist order of explanation. It starts with the idea that empirical experience presupposes the availability of determinate concepts. For apperception—awareness in the sense required for sapience, awareness that can have *cognitive* significance—is judgment: the application of concepts. Even classification of something particular as of some general kind counts as *awareness* only if the general kind one applies is a *concept*: something whose application can both serve as and stand in need of *reasons* constituted by the application of *other* concepts. When an iron pipe rusts in the rain, it is in some sense classifying its environment as being of a certain general kind, but is in no interesting sense *aware* of it. So one must already have concepts in order to be aware of anything at all.

Kant, I have said, understands *apperception*—what the transcendental unity of apperception is a unity of—which is to say judgment, in *normative* terms, which I have expressed by means of the concepts of commitment, responsibility, and endorsement. The transcendental *unity* of apperception is a normative unity: judges as such are obliged to renounce commitment to contents that are *incompatible* with their other commitments, or which have such commitments

as their consequences. For if two commitments are incompatible, each serves as a reason to give up the other. That normative unity is *transcendental* because reference to *objects*—the representational ‘of’-intentionality that Kant is concerned to show is a necessary sub-structure of inferential ‘that’-intentionality—is secured in part precisely by ‘repelling’ incompatible commitments. The judgment that A is a dog is not incompatible with the judgment that B is a fox. The judgment that A is a dog is incompatible with the judgment that A is a fox. Taking a dog-judgment to be incompatible with a fox-judgment is taking them to refer to or represent an object: the *same* object. Taking it that A is a dog does not entail that B is a mammal, but taking it that A is a dog does entail that A is a mammal. Drawing the inference is taking it that the two judgments refer to the same object. This triangulation by acknowledging incompatibilities and inferences is, in a nutshell, how the *normative* demand for a *rational* unity of apperception (judgments) makes intelligible representational purport: what it is to take or treat judgments as representing or being about objects. It follows that for concepts to perform their function in articulating the transcendental unity of apperception, the inferential and incompatibility relations they stand in to one another must be settled independently of and antecedently to our particular applications of them in judgment.

10. *Hegel and the Social Division of Labor*

Of course, this is just the point at which the pre-Kantian rationalists notoriously faced the problem of where determinate concepts come from. If they are presupposed by experiential awareness, then it seems that they cannot be thought of as derived from it, for instance by

abstraction. Once the normative apperceptive enterprise is up-and-running, further concepts may be produced or refined by various kinds of judgments (for instance, reflective ones), but concepts must always already be available for judgment, and hence apperception, to take place at all.

Empirical activity, paradigmatically apperception in the form of judgment, presupposes transcendental activity, which is the rational criticism and rectification of ones commitments, making them into a normatively coherent, unified system. Defining that normative unity requires the availability of concepts with already determinate contents (roles in reasoning).

Leibniz's appeal to innateness is not an attractive response to the resulting explanatory demand. And it would not be much improvement to punt the central issue of the institution of conceptual norms from the realm of empirical into the realm of noumenal activity. I think it is a nice question just how Kant's account deals with this issue.

As I read him, Hegel criticizes Kant on just this point. He sees Kant as having been uncharacteristically and culpably uncritical about the origin and nature of the determinate contentfulness of empirical concepts. Hegel's principal innovation is his idea that in order to follow through on Kant's fundamental insight into the essentially *normative* character of mind, meaning, and rationality, we need to recognize that *normative* statuses such as authority and responsibility are at base *social* statuses. Hegel takes it that the Enlightenment tradition was right to see normative statuses as instituted by normative attitudes. There were no such things as commitments and entitlements, responsibility and authority, before we started practically *taking* or *treating* each other *as* committed and entitled, responsible and authoritative.

Think in this connection about the example appealed to earlier, of the young one who achieves legal majority upon reaching the age of 21. The transformation in positive freedom is vast. But it is not the consequence of some magical *inner* transformation of the youth. It is *wholly* a shift in social status. *All* that changes is that others now *take* the individual to be able to commit himself, *hold* him responsible for what he does, *acknowledge* his authority so to bind himself. A Laplacian demon omniscient not only about physical occurrences, but also about Cartesian mental episodes, need not be able to discern any difference between what is going on when the subject scratches a signature on a document one day before and one day after his 21st birthday, so long as it confines its attention to what is going on under his skin and between his ears. There is indeed a difference of immense significance—but it takes place *outside* the individual. (Of course this social practice gains its point from the thought that older individuals are in general more likely to know what they are doing, what commitments and responsibilities they are undertaking, what they are authorizing, than younger ones are. But no sensible person thinks that every 22 year old understands these things better than any 20 year old.)

On this Hegelian social line, there is something importantly right about the Kant-Rousseau demarcation of the normative in terms of autonomy. We should think of each of us as bound only by the commitments we ourselves have undertaken (explicitly or implicitly). But that autonomy claim about normative *force*—that one is genuinely normatively bound only by what one has bound oneself by, commitments one has oneself endorsed—is intelligible in principle only against the background of a *social* division of labor concerning the relation between normative *force* and conceptual *content*. Here Kant's methodological individualism critically impoverishes his explanatory resources. It is an absolutely essential part of Hegel's story that we

hold *each other* responsible, acknowledge *each other's* authority. Self-regarding practical normative attitudes cannot by themselves underwrite conceptual contents that swing sufficiently free of a knower's or agent's attitudes to count as genuinely normatively *constraining* her—as articulating determinate commitments and responsibilities.

Hegel's term for the normatively articulated realm of discursive activity (Kant's "realm of freedom") is 'Geist': spirit. At its core is *language*: "Language is the Dasein of Geist," Hegel says.³ That is where concepts (which for Hegel, as for Kant, is to say, norms) have their actual, public existence. (Think of Sellars's principle that "Grasp of a concept is mastery of the use of a word.") Here is how I think the social division of conceptual labor works on Hegel's picture. It *is* up to me which counter in the game I play, which move I make, which word I use. But it is *not* then in the same sense up to me what the significance of that counter is—what other moves it precludes or makes necessary, what I have said by using that word. It *is* up to me what concept I apply in a particular judgment—whether I claim that the coin is made of copper or silver, for instance. But if I claim that it is copper, it is *not* then up to me what move I have made, what else I have committed myself to by using that term. So, for instance, I have thereby committed myself to the coin melting at 1084° C, but not at 1083° C, and to its having a density of 8920 kg/m³—in the sense that if those claims are not true then neither is the one I made. And I have made a claim that is incompatible with saying that the coin is an electrical insulator. I can bind myself by these determinate conceptual norms because they are always already there in the always already up-and-running communal linguistic practices into which I enter as a young one.

³ In the *Phenomenology of Spirit*, [A.W. Miller, (trans.), Oxford University Press] paragraph 652.

An essential part of what maintains them is the attitudes of others—in this case, of the metallurgical experts who would hold me responsible for those commitments on the basis of my performance, if the issue arose. Of course in this way the issue of the ultimate origins of concepts is only displaced, from the individual mind to the whole linguistic community—from the relatively recent to the relatively distant past. I think in fact there is a convincing story to be told about what it is for the normative “light to dawn slowly over the whole” among our hominid ancestors, but I’m not going to follow out this particular argumentative thread any further here.

11. *The Linguistic Model of Positive Freedom as Constraint by Conceptual Norms*

Instead, I want to say something about how Hegel’s social, linguistic development of Kant’s fundamental insight into the essentially normative character of our mindedness provides a model of positive freedom. One of the central issues of classical political philosophy is how to reconcile individual freedom with constraint by social, communal, or political norms. Kant’s vision of us as rational creatures opens up space for an understanding of a kind of freedom that consists in being able constrain ourselves by norms—indeed, by norms that are rational, in the sense that they are conceptual norms: norms articulating what is a reason for what. The normative conception of positive freedom then makes possible a distinctive kind of answer to the question of how the loss of individual negative freedom—freedom from constraint—inevitably involved in being subject to institutional norms could be *rationaly justified* to the individual. Even if it could be justified from the point of view of the collective—which cannot exist without

such constraints on individual behavior—it is important that it can it also be understood as rationally justifiable from the point of view of the individual herself. In the Kantian context, such a justification could in principle consist in the corresponding increase in positive freedom.

The positive expressive freedom, the freedom *to* do something, that is obtainable only by constraining oneself by the conceptual norms implicit in *discursive* social practices, speaking a public language, is a central case where such a justification evidently is available. Speaking a particular language requires complying with a daunting variety of norms, rules, and standards. The result of failure to comply with enough of them is unintelligibility. This fact can fade so far into the background as to be well-nigh invisible for our home languages, but it is an obtrusive, unpleasant, and unavoidable feature of working in a language in which one is *not* at home. The same phenomenon is manifest in texts that intentionally violate even a relatively small number of central grammatical and semantic norms, such as Gertrude Stein’s prose. But the kind of positive freedom one gets in return for constraining oneself in these multifarious ways is distinctive and remarkable.

The astonishing empirical observation with which Chomsky inaugurated contemporary linguistic theory is that almost every sentence uttered by an adult native speaker is radically *novel*. That is, not only has that speaker never heard or uttered just that sequence of words before, but neither has anyone else—ever. “Have a nice day,” may get a lot of play in the States, and “Noch eins,” in Germany, but any tolerably complex sentence is almost bound to be new.

Quotation aside, it is for instance exceptionally unlikely that anyone else has ever used a sentence chosen at random from the story I have been telling. And this is not a special property of professor-speak. Surveys of large corpora of actual utterances (collected and collated by indefatigable graduate students) have repeatedly confirmed this empirically. And it can be demonstrated on more fundamental grounds by looking at the number of sentences of, say, thirty words or less that a relatively simple grammar can construct using the extremely minimal 5000-word vocabulary of Basic English. There hasn't been time in human history for us to have used a substantial proportion of those sentences—or just the true ones—even if every human there had ever been always spoke English and did nothing but chatter incessantly. Yet I have no trouble producing, and you have no trouble understanding, a sentence that (in spite of its ordinariness) it is quite unlikely anyone has happened to use before, such as:

We shouldn't leave for the picnic until we're sure that we've packed my old wool blanket, the thermos, and all the sandwiches we made this morning.

This capacity for *radical semantic novelty* fundamentally distinguishes sapient creatures from those who do not engage in linguistic practices. Because of it we can (and do, all the time) make claims, formulate desires, and entertain goals that no-one in the history of the world has ever before so much as considered. This massive positive expressive freedom transforms the lives of sentient creatures who become sapient by constraining themselves by linguistic—which is at base to say conceptual—norms.

So in the conceptual normativity implicit in linguistic practice we have a model of a kind of constraint—loss of negative freedom—that is repaid many times over in a bonanza of positive

freedom. Anyone who was in a position to consider the trade-off rationally would consider it a once-in-a-lifetime bargain. Of course, one need not be a creature like us. As Sellars says, one always could simply *not speak*—but only at the price of having nothing to say. And non-sapient sentients are hardly in a position to weigh the pros and cons involved. But the fact remains that there *is* an argument that shows that at least *this* sort of normative constraint is rational from the point of view of the individual—that it pays off by opening up a dimension of positive expressive freedom that is a pearl without price, available in no other way. Hegel’s idea is that this case provides the model that every other social or political institution that proposes to constrain our negative freedom should be compared to and measured against. The question always is: what new kind of expressive freedom, what new kinds of life-possibilities, what new kinds of commitment, responsibility, and authority are made possible by the institution? The strategy is to use an understanding of the basic metaphysical structure of mind, meaning, and rationality as the basis for normative assessment of lives and institutions.

12. *Modality and Lawfulness*

I want to close by mentioning a topic that initially no doubt seems far-removed indeed from issues of personal autonomy and political freedom: alethic modality. Kant read Hume’s practical and theoretical philosophies as raising variants of a single question. On the side of practical reasoning, Hume asks what our warrant is for moving from descriptions of how things are to prescriptions of how they ought to be. How can we rationally justify the move from ‘is’ to ‘ought’? On the side of theoretical reasoning, Hume asks what our warrant is for moving from

descriptions of what in fact happens to characterizations of what must happen, and what could not happen. How can we rationally justify the move from descriptions of matter-of-factual regularities to formulations of necessary laws? In Kant's terminology, these are both species of 'necessity': moral or practical, and natural necessity, respectively. For him, 'necessary' (notwendig) just means "according to a *rule*". Hume's predicament is that he finds that even his best understanding of *facts* doesn't yield an understanding of *rules* governing and relating those facts, underwriting assessments of which of the things that actually happen (something we can experience) *ought* to happen (are *normatively* necessary), or *must* happen (are *naturally* necessary).

Kant's response is that Hume's predicament is not a real one, but the product of a confusion. One cannot in fact fully understand the descriptive, empirical employment of ordinary determinate concepts such as 'cat' without at least implicitly understanding also what is made explicit by the modal concepts that articulate laws. Hume thinks he can understand what it is to say that the cat is on the mat without understanding what it means to say that it is *possible* for the cat to be elsewhere, but *necessary* that it not be larger than the Earth. Kant's claim, put in contemporary terms, is that part of what one is committed to in applying any determinate concept in empirical circumstances is a distinction between counterfactual differences in circumstances that *would* and *would not* affect the truth of the judgment one is making. It would still be true that the cat is on the mat if the lighting were subtly different, but it would not be true that the cat is on the mat if the force of gravity were two orders of magnitude stronger than it in fact is. The cat could still be on the mat if the mat had twice the area it does, but not if the floor under it were not rigid—that is, not disposed to resist possible deformations of its shape.

Hume frames his question as an epistemological one, concerning the justification of our claims to know what must happen or what ought to happen based on our experience of how things in fact are. Once again, Kant offers both a semantic diagnosis of the origins of the epistemological predicament that makes this question seem urgent and difficult. And once again, he offers a semantic response that, if successful, defuses the epistemological worry.

Sellars summarizes this Kantian thought in the title of one of his essays: “Concepts as Involving Laws, and Inconceivable Without Them.” This slogan is a good place to start in thinking about Kant’s point, but in fact Sellars’s own view is subtly but importantly different from Kant’s. For Sellars, the laws determining the truth of counterfactuals involving the application of a concept are part of the content of the concept. For Kant, modal concepts make explicit not something implicit in the *content* of determinate concepts, but something implicit in their *empirical use*, in *applying* them to make empirical *judgments*. That is why the pure concepts of the understanding—what he calls ‘categories’, such as possibility and necessity—are both to be understood in terms of the forms of judgment (the table of categories derives from the table of judgments) and express synthetic, rather than analytic necessities. From Kant’s point of view, a better slogan than Sellars’s would be “The *Use* of Concepts in Empirical *Judgments* as Involving Laws and Inconceivable Without Them.”

A corresponding line of thought is to be mounted on the side of normative or practical necessity. Normative concepts make explicit commitments that are implicit in the practice of acting

intentionally, in the exercise of practical agency itself. Intentional agency is a thoroughly normative phenomenon because it, too, consists in the application of concepts, and applying concepts is undertaking commitments and responsibilities whose *content* is articulated by those concepts. For Kant, specifically moral normative vocabulary makes explicit commitments that are already implicit in the practical use of concepts to endorse maxims, ends, and plans.

My point is that Kant's response to Hume's predicament—his account of the nature and expressive role of modal and normative “pure” concepts—is not in fact as removed from his discussion of the nature of freedom as might at first have appeared. Both are rooted in and developments of his normative turn: his fundamental reconstrual of mind, meaning, and rationality in normative terms.

13. Conclusion

My aim in this essay has been to convey what in my title I call “some Kantian lessons” about what it is to have a mind, to grasp and apply meanings, to be rational. What I've been doing is not really Kant exegesis. I haven't been concerned to interpret particular bits of his text, so as to catch him expressing the views I've been attributing to him. That is an important and necessary task, and in its absence I can at most claim to have been expounding ‘kantian’ lessons, not Kant's own theory. My characterization of Kant's largest ideas and their relations to one another deserves to be controversial and is arguably tendentious. But I think that in thinking about

Kant's grandest philosophical contributions there is a standing danger of losing sight of the forest by focusing on the trees. The cost of succumbing to that danger is to fail to appreciate why Kant is so important: the conceptual sea-change he ushers in, the radically new constellation of philosophical ideas he puts in play. I think we have only really just begun the process of digesting those ideas. Though the thought sometimes tempts me, I will not in fact claim that Kant tells us nearly everything we need to know about minds, concepts, and their use and contents. But what he does tell us is so deep and significant, and ramifies into and reverberates in so many neighboring theoretical domains, that I think it does deserve to be thought of as the most important distinctively *philosophical* contribution to the multidisciplinary study of mind, meaning, and rationality. I have tried to say something here about why I think Kant is and remains for philosophers what the sea was for Swinburne: the great, grey mother of us all.

End

Panoramas of Mind and Meaning

Chapter Two :

From German Idealism to American Pragmatism—and Back

I. Kant and Hegel

Developments over the past four decades have secured Immanuel Kant's status as being for contemporary philosophers what the sea was for Swinburne: the great, gray mother of us all. And Kant mattered as much for the classical American pragmatists as he does for us today. But we look back at that sepia-toned age across an extended period during which Anglophone philosophy largely wrote Kant out of its canon. The founding ideology of Bertrand Russell and G.E. Moore, articulating the rationale and fighting faith for the rising tide of analytic philosophy, was forged in a recoil from the perceived defects of a British idealism inspired by Hegel. Mindful of the massive debt evidently and self-avowedly owed by Hegel to Kant, and putting aside neo-Kantian readings of Kant as an empiricist philosopher of science that cast him in a light they would have found more favorable, Russell and Moore diagnosed the idealist rot as having set in already with Kant. For them, and for many of their followers down through the years, the progressive current in philosophy should be seen to have run directly from Locke,

Leibniz, and Hume, to Mill and Frege, without any dangerous diversion into the oxbow of German idealism.

What did the pragmatists learn from Kant? I want to focus on two of Kant's master ideas: what I'll call his *normative turn*, and what I'll call (tendentiously but only proleptically) his *pragmatist* methodology. I think that we should still care today about these ideas—ideas which were for complicated reasons largely invisible to classical analytic philosophy. As I understand his work, Kant's most basic idea, the axis around which all his thought turns, is that what distinguishes exercises of judgment and intentional agency from the performances of merely natural creatures is that judgments and actions are subject to distinctive kinds of *normative* assessment. Judgments and actions are things we are in a distinctive sense *responsible* for. They are a kind of *commitment* we undertake. Kant understands judging and acting as applying *rules*, concepts, that determine *what* the subject becomes committed to and responsible for by applying them. Applying concepts theoretically in judgment and practically in action binds the concept user, commits her, makes her responsible, by opening her up to normative assessment according to the rules she has made herself subject to.

The responsibility one undertakes by applying a concept is a task responsibility: a commitment to *do* something. On the theoretical side, what one is committed to doing, what one becomes liable to assessment as to one's success at doing, is integrating one's judgments into a whole that exhibits a distinctive kind of unity: the synthetic unity of apperception. It is a systematic, rational unity, dynamically created and sustained by drawing inferential consequences from and finding reasons for one's judgments, and rejecting commitments incompatible with those one has undertaken. Apperceiving, the characteristically sapient sort of awareness, is discursive (that is, conceptual) awareness. For it consists in integrating judgments

into a unity structured by relations of what judgments provide reasons for and against what others. And those rational relations among judgments are determined by the rules, that is the concepts, one binds oneself by in making the judgments. Each new episode of experience, paradigmatically the making of a perceptual judgment, requires integration into, and hence transformation of the antecedent constellation of commitments. New incompatibilities can arise, which must be dealt with critically by rejecting or modifying prior commitments. New joint consequences can ensue, which must be acknowledged or rejected. The process by which the whole evolves and develops systematically is a paradigmatically rational one, structured by the rhythm of inhalation or amplification by acknowledging new commitments and extracting new consequences, and exhalation or criticism by rejecting or adjusting old commitments in the light of their rational relations to the new ones.

Kant's new normative conception of what the activity of judging consists in, of what one must be *doing* in order to be judging (a corresponding story applies to acting), puts important structural constraints on how he understands the judgeable *contents* for which one is taking responsibility in judgment. The dominant order of logical and semantic explanation of the tradition Kant inherited began with a doctrine of terms or concepts. On that base, a doctrine of judgments was erected, and then finally a doctrine of consequences or syllogisms. But the minimal unit of *responsibility* is the judgment. It is judgments, not concepts, that one can invest one's authority in, commit oneself to, by integrating them into an evolving constellation that exhibits the rational synthetic unity of apperception. Accordingly, in a radical break with his predecessors, Kant takes judgments to be the minimal units of awareness and experience. Concepts are to be understood analytically, as functions of judgment—that is, in terms of the contribution they make to judgeable contents. To be candidates for synthesis into a system

exhibiting the rational unity characteristic of apperception, judgments must stand to one another in relations of material consequence and incompatibility. So if one is to understand judging also as the application of concepts, the first question one must ask about the contents of those concepts how the use of one or another concept affects those rational relations among the judgeable contents that result. This methodological inversion is Kant's commitment to the *explanatory primacy of the propositional*. It is a methodological commitment that will be seconded by Frege, whose *Begriffsschrift* is structured by the observation that it is only judgeable contents to which pragmatic force can attach, and by Wittgenstein, who in the *Investigations* gives pride of place to sentences as the only kind of linguistic expression that can be used to make a move in a language game.

Kant's thought here, I think, is that alongside the *local* order of explanation, which looks to the contents of the particular concepts applied in judging to explain the specific possibilities of rational integration of judgeable contents containing them (their inferential grounds, consequences, and incompatibilities), there is a *global* order of explanation according to which one must understand what conceptual content *is* in terms of what judgeable contents are, and must understand that in terms of what one is *doing* in judging, in making oneself responsible for such contents. The *functionalism* about conceptual contents that consists in understanding them as functions of judgment, which is the practical expression of methodological commitment to the explanatory primacy of the propositional, is motivated by an overarching methodological *pragmatism* according to which semantics must answer to pragmatics (in a broad sense).⁴ It is the strategy of understanding discursive *content* in terms of what one is *doing* in endorsing or applying it, of approaching the notions of judgeable, and therefore conceptual content generally,

⁴ Later on (in Section V) I will suggest a somewhat narrower use of the term "methodological pragmatism".

in terms of the constraints put on it by requirement derived from the account of the *activity* of judging.

Though I have for expository reasons focused my sketch on the cognitive, theoretical side of Kant's thought, it is important to be clear that pragmatism in the sense I am attributing to Kant is *not* a matter of giving explanatory priority to the practical over the theoretical, to exercises of agency over exercises of cognition. Rather, within both the practical and the theoretical spheres, it is understanding *content* in terms of *force* (in Frege's sense): what is judged, believed, or done in terms of one must *do*, what activity one must engage in, to be judging, believing, or doing it. Kant, I am claiming, should be thought of as a pragmatist *avant la lettre* because of the way his normative theory of conceptual *activity* (theoretical and practical) shapes his account of conceptual *content* (both theoretical and practical).

I read Hegel as taking over from Kant commitment both to a normative account of conceptual doings, and to a broadly pragmatist approach to understanding the contents of our cognitive and practical commitments in terms of what we are doing in undertaking those commitments. I see him as taking an important step toward *naturalizing* the picture of conceptual norms by taking those norms to be instituted by public social recognitive practices. Further, Hegel tells a story about how the very same practice of rational integration of commitments undertaken by applying concepts that is the synthesis at once of recognized and recognizing individual subjects and of their recognitive communities is at the same time the historical process by which the norms that articulate the contents of the concepts applied are instituted, determined, and developed. He calls that on-going social, historical process "experience" (Erfahrung), and no longer sees it as taking place principally between the ears of an individual. Closer to our own time, we have seen a version of this development repeated as

Quine, in a pragmatist spirit, rejects Carnap's two-phase picture, according to which first one institutes meanings (a language) and then applies them (adopts a theory), in favor of a unitary process and practice of using expressions that must be intelligible at every stage as settling (insofar as it is settled) both what one means and how one takes things to be.

II. Classical American Pragmatism

In the broadest terms, the classical American pragmatists, Peirce, James, and Dewey, developed this German idealist tradition by completing the process of *naturalizing* it, which had begun already with Hegel. In their hands, it was to take on the shape of an empirical scientific account of us and our transactions with our environment. The sort of understanding they sought was decisively shaped by two new models of scientific explanation, codifying new forms of intelligibility characteristic of late nineteenth century science. Principal among these, of course, was Darwinian evolutionary explanations. The other form of explanation that was coming to maturity in the science of the day was statistical explanation. Pragmatism begins with a philosophy of science, pioneered by Peirce, that saw these two explanatory innovations as aspects of one conceptual revolution in science.

One dimension along which evolutionary and statistical explanations differ from those of the older mathematical physics concerns the dominant *modality* in which they are expressed. The modality of Newtonian laws is *necessity*. One explains something by showing that it is

necessitated by eternal, exceptionless, universal laws. Evolutionary and statistical explanations explain *contingent* happenings, by displaying conditions under which they can be seen to have been *probable*. Both are ways of making intelligible the contingent emergence of collective order from individual randomness.

The original subject-matter of evolutionary explanations was, of course, the process by which biological species arise and diversify. Taking his cue from the way in which statistical explanation had been generalized from its original applications in social science to provide the basis for the triumph of thermodynamics in physics, Peirce substantially generalized evolutionary-statistical forms of intelligibility in two different directions. Most important was an idea that was picked up and developed by James and above all by Dewey: the recognition that *evolution*, at the level of species, and *learning*, at the level of individuals, share a common *selectional* structure. Both can be understood as processes of *adaptation*, in which interaction with the environment preserves and reproduces (selects) some elements, while eliminating others. This insight is encapsulated in the concept of habit, and the picture of individual learning as the evolution-by-selection of a population of habits. This master idea made possible the naturalistic construal of a cognitive continuum that runs from the skillful coping of the competent predator, through the practical intelligence of primitive hominids, down to the traditional practices and common sense of civilized humans, all the way to the most sophisticated theorizing of contemporary scientists. All are seen as of a piece with, intelligible in the same general terms as, biological evolution.

The other direction in which Peirce generalized the evolutionary statistical selectional model of explanation was to *inorganic* nature. What those older scientific naturalists, for whom the paradigm of scientific understanding was Newtonian physics rather than Darwinian biology,

had taken to be eternal, immutable, necessary, universal laws of nature, Peirce now sees as themselves in the largest sense “habits” of the universe—a kind of order that has arisen contingently, but ultimately statistically explicable, by a selectional-adaptational process operating on a population of such regularities, which in turn provides the dynamic habitat to which all must collectively adapt. There is no guarantee that any such accommodation will succeed permanently. As with habits learned by individuals, some of the lawlike regularities may prove more robust and others more fragile. The older picture of laws shows up as at best only approximately true, an idealization extrapolating a situation that actuality approaches at most asymptotically.⁵ The naturalism of the classical American pragmatists was shaped by the new sort of nature they had been taught about by the best science of their times—a nature viewed through the lens of the new forms of statistical and selectional explanation.

The pragmatists’ new form of *naturalism* was coupled with a new form of *empiricism*. The experimental scientific method is seen as just the explicit, principled distillation of the selectional learning process that is the practical form common to intelligent creatures at all stages of development. Dewey’s term for that process, in all its varieties, is ‘experience’—the axial concept of such central works as *Experience and Nature* and *Art as Experience*. (So central is the concept to Dewey’s thought that sometimes in reading these works it is difficult to overcome the impression that he is, as Rorty once put it, “using the term ‘experience’ just as an incantatory device to blur every conceivable distinction.”) Experience in this sense is not the ignition of some internal Cartesian light—the occurrence of a self-intimating event of pure awareness, transparent and incorrigible to the subject of the experience. Experience is *work*: the application of force through distance. It is something *done* rather than something that merely *happens*—a

⁵ James endorses this Peircean idea in Lecture II of *Pragmatism*.

process, engaging in a practice, the exercise of abilities, rather than an episode. It is experience, not in the sense of *Erlebnis* (or *Empfindung*), but of Hegel's *Erfahrung*. It is the decidedly non-Cartesian sense of 'experience' in which a want-ad can specify "No experience necessary," without intending thereby to invite applications from zombies. Earlier empiricists had thought of experience as the occurrence of conscious episodes that provide the raw materials for learning, via processes such as association, comparison, and abstraction. For the pragmatists, experience is not an input to the learning process. It just *is* learning: the process of perception and performance, followed by perception and assessment of the results of the performance, and then further performance, exhibiting the iterative, adaptive, conditional-branching structure of a Test-Operate-Test-Exit loop. The result of experience is not best thought of as the possession of items of *knowledge*, but as a kind of practical *understanding*, a kind of adaptive attunement to the environment, the development of habits apt for successful coping with contingencies. It is knowing how rather than knowing that.

Ontological naturalism and epistemological empiricism are both encouraged by the idea that the rise of modern science, the most successful social institution of the past three hundred years, can teach philosophers the most important lessons both about how things are and how we can understand them. But from the beginning they have typically stood in significant tension with one another. The furniture of Newton's natural world does not include Locke's mind. And Hume can find nothing in experience by which we could come to know or understand laws such as Newton's as having the necessity that distinguishes laws from mere regularities. Nor is this tension a characteristic only of Enlightenment naturalism and empiricism. It equally afflicts the twentieth-century versions. The two principal wings of the Vienna Circle, which Carnap struggled heroically to keep from flying off in different directions, were distinguished precisely

by their answers to the question: when empiricism and naturalism conflict, which should be relaxed or given up? Schlick urged the preeminence of empiricism, while Neurath was committed to the priority of naturalism. Quine never fully reconciled his (logical) empiricist hostility to modality with his naturalist privileging of the deliverances of science.

The classical pragmatist versions of naturalism and empiricism, though, fit together much better than the versions that preceded and succeeded them. Far from being in tension, they complement and mutually support one another. Both the world and our knowledge of it are construed on a single model: as mutable, contingent products of statistical selectional-adaptational processes that allow order to pop to the surface and float in a sea of random variability. Both nature and experience are to be understood in terms of the processes by which relatively stable constellations of habits arise and sustain themselves through their interactions with an environment that includes a population of competing habits. There is no problem in principle in finding a place for experience construed as learning in nature construed as evolving. Nor is there any analog of the traditional complementary problem of understanding how experience construed as the dynamic evolution of habits can give its subjects access to the modally robust habits of the things those knowers-and-agents interact with, adapt, and adapt to. The pragmatist forms of naturalism and empiricism are two sides of one coin.

The pragmatists' conception of experience is recognizably a naturalized version of the rational process of critically winnowing and actively extrapolating commitments, according to the material incompatibility and consequence relations they stand in to one another, that Kant describes as producing and exhibiting the distinctive synthetic unity of apperception. For that developmental process, too, is selectional (though not statistical). Some commitments (theoretical and practical) thrive and persist, in concert with their fellows, while others are

modified or rejected as unable to flourish in that environment. It might be thought fanciful to focus on this common structure in light of the substantial difference between the conceptions: Kant's process is structured by *rational*, conceptual relations of incompatibility and consequence, while the pragmatists' version is structured by *natural*, causal relations of incompatibility and consequence.

But the pragmatists would disagree. For they introduce not only a new conception of experience, but also a new conception of reason. They understand the rationality of the theoretical physicist as continuous with the intelligence of the culturally primitive hunter and the skill of the non-human predator. The grooming and development of discursive cognitive and practical commitments is a learning process of a piece and sharing a structure with the achievement of practical attunement to an environment and the acquisition of habits successful in that environment that in one form or another is a part of the natural history of all sentient organisms. Reason and intelligence in this sense can be seen (albeit in an inflexible and unlearned form) already in the maintenance of an equilibrium by that emblem of the industrial revolution: the fly-wheel governor. The nature of the pragmatists is through and through a rational nature—not just the part of it that is intelligible as experience.

III. Fundamental Pragmatism

The more specific strategy by which the classical American pragmatists sought to naturalize the concept of experience—to demystify and domesticate it, to disentangle it from two centuries of Cartesian encumbrances—is what I will call *fundamental* pragmatism. This is the idea that

one should understand knowing *that* as a kind of knowing *how* (to put it in Rylean terms). That is, believing *that* things are thus-and-so is to be understood in terms of practical abilities to *do* something. Dewey, in particular, saw the whole philosophical tradition down to his time as permeated by a kind of platonism or intellectualism that saw a rule or principle, something that is or could be made conceptually or propositionally explicit, behind every bit of skillful practice. He contrasted that approach with the contrary pragmatist approach, which emphasizes the implicit context of practices and practical abilities that forms the necessary background against which alone states and performances are intelligible as explicitly contentful believings and judgments. In this reversal of the traditional order of explanation, Dewey is joined by the Heidegger of *Being and Time*, with his project of understanding *Vorhandenheit* as a precipitate of the more ‘primordial’ *Zuhandenheit*, and by the later Wittgenstein. All three thinkers are downstream from Kant’s fundamental insight about the normative character of cognition and agency, and share a commitment to the explanatory priority of norms implicit as proprieties of practice to norms explicit as rules or principles.

I mean the rubric “fundamental pragmatism” to be a relatively loose and elastic description, whose parameters can be adjusted or interpreted so as to fit the methodology of many thinkers, who might differ in many other ways. It is supposed, for instance, to include both the order of explanation that lead Quine to criticize “myth of the museum” in thinking about meaning and that Sellars employs in criticizing the “myth of the given” in thinking about sensory experience. It depends on a contrast, which may be filled-in in different ways, between something on the implicit, know-how, skill, practical ability, practice side and something on the explicit, conceptual, rule, principle, representation side. So we might distinguish between two grades of intentionality: practical and discursive. Practical intentionality is the kind of attunement to their

environment that intelligent nonlinguistic animals display—the way they can practically take or treat things as prey or predator, food, sexual partner or rival and cope with them accordingly. Discursive intentionality is using concepts in judgment and intentional action, being able explicitly to take things to be thus-and-so, to entertain and evaluate propositions, formulate rules and principles. The fundamental pragmatist aspiration is to be able to exhibit discursive intentionality as a distinctive kind of practical intentionality. This project can take a strong reductionist form. For instance, what I have elsewhere⁶ called the “pragmatist version of artificial intelligence” claims that there is a set of practices or abilities that are non-discursive, in the sense that each of them can be engaged in or exercised by nondiscursive creatures, and yet which can be algorithmically elaborated into the discursive capacity to use concepts and speak an autonomous language. But fundamental pragmatism need not take such a strong, reductive form. One might claim, more modestly, that discursive activity, from everyday thought to the cogitations of the theoretical physicist, is a species of practical intentionality (or a determination of that determinable), and indeed, one that is intelligible as having developed out of nondiscursive practical intentionality, while still maintaining that it is a wholly distinctive variety.

Fundamental pragmatism in this sense gives a distinctive shape to the naturalism of the classical American pragmatists. For that methodological commitment ensures that their naturalism is in the first instance a naturalism concerning the *subjects* of discursive understanding and agency. When we think today about naturalism, we tend to think of it first as a thesis about the *objects* represented by different potentially puzzling kinds of concepts: semantic, normative, probabilistic concepts, and so on. The question is how to see what those

⁶ In Chapter 3 of *Between Saying and Doing* [Oxford University Press, 2008].

concepts represent as part of the natural world, as conceived by fundamental physics, or some special sciences, or even just by unproblematic empirical descriptive concepts. By contrast to this object naturalism, the American pragmatists were subject naturalists.⁷ Fundamental pragmatism counsels looking first to what discursive subjects are *doing*, to the abilities they exercise, the practices they engage in. If a naturalistic story can be told about *that*, it might well be that no questions remain that should trouble the naturalist. One of the points of the toy Sprachspeile that the later Wittgenstein constructs seems to be a fundamental pragmatist, subject naturalist one—which the distinction between subject and object naturalism shows to be entirely compatible with the claim he makes already in the *Tractatus* and never relinquishes, that “philosophy is not one of the natural sciences.” Not everything we think or say need be understood as representing the world as being some way. And if it is, fundamental pragmatism invites us to understand representation in terms of what discursive subjects must *do* in order to count thereby as representing, as taking or treating some state, episode, or performance *as* a representation *of* something. For representational content is explicit—believing that things are thus-and-so. And that is to make sense of in terms of what is implicit in what the subjects do in virtue of which it is correct to say of them that they are believing that. Fundamental pragmatism is opposed to a *representationalist* order of explanation: one that begins with a notion of representational content, and appeals to that to make sense of what it is knowing and acting subjects do. That is not to say that pragmatists in this sense can have no truck at all with the

7 This is Huw Price’s terminology in “Naturalism without representationalism” in David Macarthur and Mario de Caro (eds), *Naturalism in Question* (Harvard University Press, 2004), 71—88, and (with David Macarthur) “Pragmatism, quasi-realism and the global challenge” In Cheryl Misak, ed., *The New Pragmatists* (OUP, 2007), 91—120. The other essays in his *Naturalism Without Mirrors* [Oxford University Press, 2009] can also be consulted with profit in this connection.

concept of representation. It is to say at most that talk of representation should come at the end of the story, not the beginning.

Once a contrast between skillful *practice* and explicit *representation* has been put in place and the issue raised of their relative explanatory priority in the context of different enterprises, the question of the relation between fundamental pragmatism and cognitive science arises. For cognitive science had as something like its original charter distinguishing its approach from that of behaviorism by its realization of the explanatory power precisely of appealing to *representations* to explain various practical cognitive abilities. Thinking about the fundamental pragmatism motivating Heidegger in setting out the project of *Being and Time*, Dreyfus drew the conclusion that the methodology of cognitive science is incompatible with the insights of that pragmatism. Is he right?

Here I think the beginning of wisdom is the realization that it makes a *big* difference whether we are talking about representations, rules, and explicitness at the *personal* level, or at the *sub-personal* level. This is in part a matter of whether one construes the *rules* the platonist invokes to articulate proprieties of practice as being *followed by* the one whose practice is in question (which would be at the personal level). Cognitive science, by contrast, postulates *sub-personal* representations, whose role is in *causal* explanations of various capacities. The sense in which they *guide* the practice is *causal*, not in the first instance *normative*. It is not at all clear that there is (or at any rate needs to be) a clash between fundamental pragmatism at the personal level and cognitive science's invocation of representations at the sub-personal level—as Dreyfus at least sometimes seems to think there is. Here one important issue is what one means by 'explicit' when fundamental pragmatism is articulated in terms of the implicit in practice vs. explicit in principle, rule, or representation form. Representations of rules are crucial for one to count as

following a rule (as Sellars insists). In that context, representation can be thought of as the form of explicitness. But it is *not* a good idea to explicate explicitness in terms of representation if one is thinking of representation in the sense that is appropriate to the *sub*-personal level. Here the notion of specifically propositional representations is key. It is open to the pragmatist to claim (with Davidson and the author of *Making It Explicit*) that *nothing* at the sub-personal level deserves to count as *propositionally* contentful in the sense that *personal* level representations can be propositionally contentful. Belief on such a view is not a sub-personal level concept.

In order to understand the relations between fundamental pragmatism and the representational approach of cognitive science we should distinguish three levels:

- a) Sub-personal representations,
- b) Practical abilities (practices) that are cognitive in some broad sense,
- c) Personal level representations.

(c) is the *explicit* properly propositional level, at which rules and principles are formulated that can express what is *implicit* at level (b). Level (b) is *practical* intentionality, and level (c) is *discursive* intentionality. Level (a) *causally* explains level (b)—and a lot of cognitive science is concerned with how this can be done in detail. The fundamental pragmatist claim is that level (c) is to be understood, explained, or explicated in terms of level (b). Cognitive science is in the business of postulating inner sub-personal representations in order to explain various kinds of skillful practice or ability. Dreyfus seems to think that approach is incompatible with the sort of fundamental pragmatism that the early Heidegger (and the later Wittgenstein) endorse. But such

a view is mistaken. What that pragmatism is incompatible with is seeking to explain (b) in terms of (c), *not* (b) in terms of (a).⁸

IV. Instrumental Pragmatism

One of the variant rough formulations I offered of the methodological commitment I have called “fundamental pragmatism” is to think about norms explicitly represented in the form of rules or principles only in the context of a prior understanding of norms implicit in practice. This characterization has the advantage of placing fundamental pragmatism in the context of the Kantian normative turn, as I have claimed it should be when we think about the classical American pragmatists. The master argument for fundamental pragmatism about the normative dimension of intentionality is a regress argument familiar from the later Wittgenstein. In a nutshell, it is that the very idea of norms explicitly represented as rules or principles presupposes that of norms implicit in practices. For applying a rule is itself something that can be done correctly or incorrectly. If we can only understand that normative assessment in turn as a matter of applying some other rule (what Wittgenstein calls an “interpretation” [Deutung]), then we are embarked on a fruitless regress. This, too, is a point that Kant had already appreciated, as an integral part of his ground-breaking normative construal of concepts as rules (for judging):

⁸ In *Between Saying and Doing* [Oxford University Press, 2008] I explore the significance of the choice of the vocabulary used to specify the practices-or-abilities appealed to at level (b). This is, it seems to me, equally significant for the two enterprises, both the one that seeks to explain them and the one that seeks to use them to explain something else.

If understanding in general is to be viewed as the faculty of rules, judgment will be the faculty of subsuming under rules; that is, of distinguishing whether something does or does not stand under a given rule (*casus datae legis*). General logic contains and can contain no rules for judgment...If it sought to give general instructions how we are to subsume under these rules, that is, to distinguish whether something does or does not come under them, that could only be by means of another rule. This in turn, for the very reason that it is a rule, again demands guidance from judgment. And thus it appears that, though understanding is capable of being instructed, and of being equipped with rules, judgment is a peculiar talent which can be practised only, and cannot be taught.⁹

The normative fundamental pragmatism of the classical American pragmatists joins cognitive science in rejecting the descriptive-dispositional behaviorism of Watson, Skinner, and Ryle. But it does so for different reasons: because of the failure of the latter group to appreciate the essentially normative character of the practical intentionality that forms the background of discursive intentionality, rather than because of their hostility to the postulation of inner representations. Wittgenstein has been called a ‘behaviorist’, in part because of his antipathy towards some kinds of explanations that appeal to inner representations. A principal danger of talking this way is that it invites overlooking his emphasis not only on the *social*, but especially on the *normative* character of the practical intentionality in the context of which he urges us to think about discursive intentionality. In this regard, Wittgenstein belongs in a box with the

⁹ *Critique of Pure Reason*, A132/B171. I discuss this regress argument further in Chapter One of *Making It Explicit* [Harvard University Press, 1994], in the context of arguments against the twin dangers of regulism and regularism about discursive norms.

classical American pragmatists, and with Kant, not with the reductive descriptive-dispositional behaviorists.

But how, exactly, do the classical American pragmatists understand the basic kind of normativity implicit in practical intentionality: the kind of skillful know-how, as a species of which we are to understand discursive intentionality and its distinctive kind of normativity? I think it is not so easy to extract a clear answer to this question, even from Dewey, who has the most sophisticated approach to it. It is clear that in the most general terms the response takes the form of an appeal to the selectional-adaptional structure common to learning and evolution. The norms characteristic of the kind of practical intentionality in terms of which we are to understand discursive intentionality are immanent to and elaborated within the development of courses of experience that display this structure. In our own time, we have examples of how to make an account along these lines work. A splendid instance is Ruth Millikan's sophisticated and nuanced construction of norms in the form of Proper Functions, defined by modal counterfactual claims about selectional processes shaping reproductive families of traits.¹⁰ (Millikan, a Sellars student, self-consciously takes her inspiration from Charles Morris, to whom her book is dedicated. Morris was a student of George Herbert Meade, who was in turn a student of James and a colleague of Dewey's.) I think there is every reason to believe that all of the classical American pragmatists (as well as the successors just mentioned) would have welcomed and embraced her careful working-out of their underlying idea. But of course, that detailed account was not available to them. In its absence, they often enough fall into formulations that have,

¹⁰ *Language, Thought, and Other Biological Categories* [MIT Press, 1984]. The basic connection between selectional processes and alethic modal counterfactuals is indicated already by Elliot Sober's distinction between traits that are selected vs. traits that are selected *for* (*The Nature of Selection* [MIT, Bradford Press, 1984]). Millikan takes the thought much farther.

from the very beginnings of the movement, led critics to attribute to the pragmatists commitment to quite a different, though not wholly unrelated, theory and to take it as the very core of the pragmatist approach.

I have in mind what is expressed by F.C.S. Schiller's slogan "The truth is what works." This is what Dewey calls "the instrumental theory" or "instrumentalism." He endorses it in such passages as these:

What should it mean upon the instrumental theory to accept some view or idea as true upon social credit? Clearly that such an acceptance itself works.¹¹

What the experimentalist means is that the effective working of an idea and its truth are one and the same thing—this working being neither the cause nor the evidence of truth but its nature...¹²

Naturally, the pragmatist claims his theory to be true in the pragmatic sense of truth: it works, it clears up difficulties, removes obscurities, puts individuals into more experimental, less dogmatic, and less arbitrarily sceptical relations to life; aligns philosophic with scientific method; does away with self-made problems of epistemology; clarifies and reorganizes logical theory, etc. He is quite content to have the truth of his theory consist in its working in these various ways, and to leave to the intellectualist the proud possession of a static, unanalyzable, unverifiable, unworking property.¹³

James says such things as:

11 "A Reply to Professor Royce's Critique of Instrumentalism" *Middle Works* Vol. 7, p. 75.

12 "The Intellectualist Criterion of Truth" *Middle Works* Vol. 4 p. 69.

13 "A Short Catechism Concerning Truth" *Middle Works* Vol 6 p. 10.

We here assume Japan to exist without ever having been there, because it WORKS to do so, everything we know conspiring with the belief, and nothing interfering...¹⁴

On pragmatistic principles, if the hypothesis of God works satisfactorily in the widest sense of the word, it is true.¹⁵

Semantic norms are understood in instrumental terms, in terms of utility. Truth-evaluable states such as beliefs are thought of on the model of tools, which can be more or less apt or useful, in concert with others that are available in a concrete situation, relative to some desired end or purpose. Taking my cue from Dewey's terminology, I'll call this approach "instrumental pragmatism" about semantic norms. There are two principal points about which it is important to be clear in thinking about the instrumental strain in classical American pragmatism. First, it should be understood as at base a theory of *meaning*, not a theory of *truth*. The pragmatists did themselves no favors by pitching it in the latter way. The general idea is the fundamental pragmatist one: that the contentfulness of intentional states such as belief should be understood in terms of the contribution they make to what the believers *do*. The new element is that the doing is thought of as purposive, as aimed at some kind of end, at the satisfaction of some desire or need. Identifying success in the doing with the *truth* of the items to be thought of as contentful in virtue of their role in that process is a further, optional move. It threatens to overshadow the underlying account of meaning and content.

14 *Pragmatism* Lecture VI.

15 *Pragmatism* Lecture VIII.

The second point is that that theory of the contentfulness of intentional states is a *functionalist* account. Instrumental pragmatism is a comprehensive holist functionalism about the content of states such as beliefs, desires, and intentions. It is comprehensive in that the functional systems considered comprise the organism and its whole environment. The role in such a functional system that determines the contents of states and performances caught up in it is a role in the process by which the system develops, through cycles of perception, thought, intervention transforming the environment, and perception of the results of that transaction. This is role in a course of *experience*, in what is very much a naturalized version of Hegel's sense of that term—a notion of experience that was in turn already a somewhat naturalized descendant of Kant's process of synthesis (by rational amplification, criticism, and justification) of something that exhibits the structure and unity of apperception. Processes of this sort involve felt dissatisfactions with the situation as it is at one moment, attempts to diagnose the nature of those felt dissatisfactions and to address and remove them, a process that, when all goes well, is at once the clarification of the dissatisfaction and its dissolution—the transformation of the old situation into a new one that is dissatisfying in some other way. That Kantian ancestry is particularly evident in some formulations of instrumental pragmatism. Here is one by James:

A new opinion counts as 'true' just in proportion as it gratifies the individual's desire to assimilate the novel in his experience to his beliefs in stock. It must both lean on old truth and grasp new fact; and its success... in doing this, is a matter for the individual's appreciation. When old truth grows, then, by new truth's addition, it is for subjective reasons. We are in the process and obey the reasons. That new idea is truest which performs most felicitously its function of satisfying our double urgency. It makes itself true, gets itself classed as true, by the way it works; grafting itself then upon the ancient

body of truth, which thus grows much as a tree grows by the activity of a new layer of cambium.¹⁶

Friction with stubborn reality is an integral component in this sort of process. That is the objective element in James's "double urgency." Feedback-regulated practices are 'thick', in the sense of essentially involving objects, events, and worldly states of affairs. Bits of the world are *incorporated* in such practices, in the exercise of such abilities. In this regard they contrast with words and sentences, considered merely as sign-designs or items in the natural world, which are 'thin' in that they can be specified independently of a specification of the objects or states of affairs they refer to or represent. For you cannot say what, for instance, the practice of attaching two boards with a hammer and nails is without referring to the boards, nails, and hammer. Dewey thinks of the thickness of pragmatist semantics as one of its cardinal advantages over its more traditional thin rivals. If one focuses on success as the measure of truth, rather than on functionalism about meaning, and further fails to appreciate that the functional system being considered is capacious enough to include the environment being acted on and in as well as the organism transacting with it, one will misunderstand instrumental pragmatism as a radically subjectivist view, according to which all that matters for truth is subjective feelings, and objective constraint vanishes. This is what I call "vulgar" pragmatism. James complains about this flat-footed, reductive reading already in *Pragmatism*:

Schiller says the true is that which 'works.' Thereupon he is treated as one who limits verification to the lowest material utilities. Dewey says truth is what gives

16 *Pragmatism* Lecture II.

'satisfaction.' He is treated as one who believes in calling everything true which, if it were true, would be pleasant.¹⁷

And he spends most of *The Meaning of Truth* (a book that on my interpretation would better have been called *The Truth About Meaning*) rebutting that reading. Here is Dewey responding to this subjectivizing reading:

Pupil: Objection Nine. Still the pragmatic criterion, being satisfactory working, is purely personal and subjective. Whatever works so as to please me is true. Either this is your result (in which case your reference to social relations only denotes at bottom a number of purely subjectivistic satisfactions) or else you unconsciously assume an intellectual department of our nature that has to be satisfied; and whose satisfaction is truth. Thereby you admit the intellectualistic criterion.

Teacher: Reply. We seem to have got back to our starting-point, the nature of satisfaction. The intellectualist seems to think that because the pragmatist insists upon the factor of human want, purpose, and realization in the making and testing of judgments, the impersonal factor is therefore denied. But what the pragmatist does is to insist that the human factor must work itself out in cooperation with the environmental factor, and that their coadaptation is both "correspondence" and "satisfaction." As long as the human factor is ignored and denied, or is regarded as merely psychological (whatever, once more, that means), this human factor will assert itself in irresponsible ways. So long as, particularly in philosophy, a flagrantly unchastened pragmatism reigns, we shall find, as at present, the most ambitious intellectualistic systems accepted simply because of the personal comfort they yield those who contrive and accept them.

17 *Pragmatism* Lecture VI

Once recognize the human factor, and pragmatism is at hand to insist that the believer must accept the full consequences of his beliefs, and that his beliefs must be tried out, through acting upon them, to discover what is their meaning or consequence.¹⁸

The possible misunderstanding is, I think, actual in the reference to "our needs" as a criterion of the correctness of truth of an idea or plan. According to the essays, it is the needs of a situation which are determinative. They evoke thought and the need of knowing, and it is only within the situation that the identification of the needs with a self occurs; and it is only by reflection upon the place of the agent in the encompassing situation that the nature of his needs can be determined. In fact, the actual occurrence of a disturbed, incomplete, and needy situation indicates that my present need is precisely to investigate, to explore, to hunt, to pull apart things now tied together, to project, to plan, to invent, and then to test the outcome by seeing how it works as a method of dealing with hard facts. One source of the demand, in short, for reference to experience as the encompassing universe of discourse is to keep us from taking such terms as "self," "my," "need," "satisfaction," etc., as terms whose meanings can be accepted and proved either by themselves or by even the most extensive dialectic reference to other terms.¹⁹

Here Dewey emphasizes not only the importance of the functionalism being comprehensive in considering a developing functional system that encompasses environment as well as the striving knower-agent, but also the holism about content that such a functionalism entails. There is no

18 "A Short Catechism Concerning Truth" *Middle Works* Vol. 6 p. 11.

19 Introduction to *Essays in Experimental Logic*, *Middle Works* Vol. 10, p. 364.

antecedently specifiable determinate content that a belief has, apart from its fellows and in advance of participating in a cycle of experience, which can then be judged *true* by pragmatist standards should the cycle conclude successfully. Rather, the belief is intelligible as having the content it does only insofar as it acquires that content by playing the role it does, along with its concomitant states, in the transactions between the believer and her world.

What goes for beliefs goes also for desires. Ends and purposes themselves are to be understood as having their content as a matter of their role in this overall system and its developmental processes. What might start out as a vague dissatisfaction itself can be clarified during the course of experience in which finding out how things are and finding out what one wants are two aspects of one process. The satisfaction of needs and wants, the achievement of goals and purposes, is the source of normativity on the instrumental construal; doing that is “working”. But what they are is (like the contents of the beliefs we are working with) itself part of what is to be determined in the course of inquiry—‘determined’ both in the sense of being made more definite and in the sense of being discovered. The former shows up from a prospective perspective, and the latter from a retrospective perspective. Dewey expended a great deal of effort in the dual process of trying to make clear and get clear himself about how the norms and standards and what they are norms and standards for assessing jointly develop in the course of experience. I cannot say that it seems to me that he succeeded very well at either task. But I do think that there is an important thought that he was after: an essentially historical perspectival structure of discursive normativity articulating a conception of determinate conceptual content that I see as also the key to understanding Hegel’s conception of experience. I have myself expended considerable effort in the dual task of trying to make that conception

clear and get clear about it myself—with what success remains to be seen. I am not going to rehearse those efforts here.²⁰

Dewey's and James's instrumentalism arises as one (optional) way of elaborating what is often called "Peirce's Principle": the meaning of a claim is the difference that adopting it would make to what one does. In fact, as I argue in Chapter Eight of *Making It Explicit*, one can get a lot more from this principle if one bifurcates it by keeping separate sets of books on the difference it makes to what one *tries* to do and difference it makes to what one *succeeds* in doing. The first of these gives one a practical difference *de dicto*. The second gives one a practical difference *de re*. Further articulating Peirce's Principle to as to take account of the intimate *social* perspectival relations between these two sorts of practical consequence—the sense in which they are two sides of one coin—allows a much more fine-grained account of conceptual content than the classical American pragmatists managed to formulate. But that, too, is a story for another occasion. The combination of the distinction of *historical* points of view between prospective (determining as clarifying) and retrospective (determining as discovering) perspectives, and *social* points of view between attributing (*de dicto*) and acknowledging (*de re*) commitments is one of my principal suggestions for how to move forward with the ideas of the classical pragmatists.

20 I talk about this structure in Chapter Three of *Reason in Philosophy* [Harvard University Press, 2009], and in Chapter Seven of *Tales of the Mighty Dead* [Harvard University Press, 2002]. It is the principal topic of my big work-in-progress on Hegel, *A Spirit of Trust*.

V. The Linguistic Turn

When classical American pragmatism is looked back upon from the perspective of the analytic movement that dominated Anglophone philosophy for at least the last half of the twentieth century, it can easily appear that a decisive wrong turn was taken after Peirce. The pragmatist founder-member was principally concerned to advance the philosophical understanding of modern logic, symbolic and natural languages, and the natural sciences—a constellation of topics that remained at the center of the analytic tradition. In his logic of relations Peirce independently achieved the bonanza of expressive power that Russell saw in Frege’s logic. But what did his successor pragmatists make of that achievement? Particularly in contrast to what Russell made of Frege, it would seem from a later vantage point that an opportunity was missed. James had little interest in logic and wrote almost nothing about it—in striking contrast to his Hegelian colleague Josiah Royce, who saw in the algebraic constructions of Alfred Bray Kempe (whom he had learned about from Peirce) a tool with which he hoped to solve the riddle of how to elaborate spatio-temporal relations from a purely conceptual basis.²¹ The logic Dewey wrote his late, important book about was unrecognizable as such to those of his

²¹ See Bruce Kuklick’s discussion of this fascinating late project in his *Josiah Royce: An Intellectual Biography* [Hackett, 1985].

readers in 1938 whose paradigm of logic was to be found in the works of Frege, Russell, and Carnap. The only pragmatist whose concern with logic matched and was recognizable as continuing that tradition was the homegrown neo-Kantian C. I. Lewis, the founder of twentieth century modal logic, who saw his own work as an attempt to synthesize the approaches of his teachers James and Royce, and in turn passed on pragmatist ideas to his students, Quine and Goodman.

Again, although James was surely the by far the best writer among the classical triumvirate, his philosophical interests focused on experience, rather than language. Dewey did write a lot about language—what he called the “tool of tools.”²² He has many good things to say about the relations between meaning and use (particularly in Chapter 5 of *Experience and Nature*). But he, too, would not be recognizable to later philosophers of language as one of their number. As for science, it is not the case that James and Dewey did not care about science and the philosophy of science. But where Peirce focused on the natural sciences, James’s contributions lay on the side of psychology, and Dewey’s main interests were in the social sciences.

By “the linguistic turn” here I mean putting language at the center of philosophical concerns, and understanding philosophical problems to begin with in terms of the language one uses in formulating them. But there is a more specific significance one can take language to have. By ‘lingualism’ (compare: ‘rationalism’)—admittedly an unlovely term—is shall mean commitment

²² *Experience and Nature, Later Works* Vol. 1, p. 134.

to understanding *conceptual* capacities (discursiveness in general) in terms of *linguistic* capacities. Dummett epitomizes a strong version of this order of explanation:

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.²³

A weaker version of lingualism claims only that language is a *necessary* condition of discursiveness, not that it is a sufficient condition that can at least in principle be made intelligible independently of talk about discursive commitments.

It would be a mistake to conclude that the pragmatists after Peirce missed the linguistic turn. In fact, Dewey at least is clearly a (weak) lingualist about the discursive. What the pragmatists did was develop these thoughts within the context of a different approach to understanding the crucial phenomenon of language—one that was complementary to that of the analytic tradition. The Frege-Russell-Carnap approach to language takes as its paradigm artificial, formal, logistical languages articulated by explicit rules. The American pragmatists, like their fellow fundamental pragmatists the Heidegger of *Being and Time* and the Wittgenstein of the *Philosophical Investigations*, address natural languages, which they think of anthropologically, as aspects of the natural history of a certain kind of being. Their focus to begin with is not on *meaning*, but on *use*: on discursive practices, skills, and abilities, on what one must be able to *do* in order to count as saying or thinking *that* things are thus-and-so.

We can think of these two approaches as distinguished by their preferred order of explanation. The question is: which comes first, semantics (the theory of meaning) or

²³ Frege's *Philosophy of Language* [ref.] p. 361.

pragmatics (the theory of use)? The logistical tradition begins with semantics: stipulating the association of some kind of semantic interpretants (paradigmatically, extensions) with basic expressions and deriving associations for more complex ones, or stipulating basic rules of derivation and then seeing what consequence relation they jointly determine. The question of how it is appropriate to *use* expressions governed by those rules is then deferred to a subsequent pragmatic theory, to which this current of thought has not traditionally devoted a great deal of attention. By contrast, the pragmatist tradition begins with pragmatics: an account precisely of how it is appropriate to use expressions. It is true that the pragmatists, also have not traditionally given a lot of attention to the specifics of the semantics that goes with such a pragmatics.

But I think we can see two principles that govern fundamental pragmatists' understanding of the relation between pragmatics and semantics. They express complementary aspects of the sense of the pragmatism in the philosophy of language that consists in insisting that semantics must answer to pragmatics. First is what I shall call "methodological pragmatism." This is the principle that the *point* of associating meanings, extensions, contents, or other semantic interpretants with linguistic expressions is to codify (express explicitly) proprieties of *use*. I think we can discern commitment to this methodological principle even in a semantic nihilist such as the later Wittgenstein. For one thing he means by saying that language is a motley is that so many and so various are the uses of any expression that there are no realistic prospects of systematizing them by associating some underlying meaning, on the basis of which one hopes then uniformly to derive the various uses (say, by one rule for declarative uses, and another for imperative ones, another for hypothetical, and so on). If the variety of uses is open-ended and

unsurveyable, then there is no prospect for semantic theorizing in philosophy, precisely because the only point of such theorizing would be systematizing those proprieties of use.

The second principle governing the pragmatists' understanding of the sense in which semantics should answer to pragmatics is what I shall call "semantic pragmatism." This is the principle that in a natural language, all there is to effect the association of meanings, contents, extensions, rules, or other semantic interpretants with linguistic expressions is the way those expressions are *used* by the linguistic practitioners themselves. *Formal* semantics for *artificial* languages can content itself with the explicit stipulation of such rules or associations of meanings, by the semantic theorist working in a semantic metalanguage. *Philosophical* semantics for *natural* languages is obliged to say what it is about the practices the users of those expressions engage in or the abilities they exercise, in virtue of which they should be understood as *governed* by those rules, or as *conferring* those meanings. Semantic pragmatism is a kind of use-functionalism about meaning (the classical American pragmatists being *comprehensive* functionalists, in the sense I have given that qualification). Again, given his practice, I think commitment to such a principle can be attributed even to such a semantic pessimist as the later Wittgenstein, precisely in virtue of his criticism of various traditional ways of thinking about meaning or content for their failure to live up to this requirement. And that sort of strategy is equally evident in Dewey's criticisms of traditional intellectualist and mentalistic conceptions.

The combination of methodological and semantic pragmatism, the two senses in which semantics can be taken to answer to pragmatics, broadly construed, might be called "linguistic pragmatism." It is one natural way of applying fundamental pragmatism to systematic theorizing

about language. One of the clearest and most emphatic proponents of that conjunctive doctrine among recent philosophers is Dummett—though of course he does not associate it with pragmatism.

Quine carries forward this general pragmatist tradition in the philosophy of language when he criticizes Carnap's two-stage picture of language, according to which first meanings are stipulated, and only subsequently are theories formulated to determine which of the sentences with those meanings are true. That division of labor makes sense for artificial languages. But to understand natural languages we have to understand how the *one* thing we do, *use* the language, can serve at once to settle the meanings of our expressions *and* determine which of them we take to be true. Linguistic practice is not illuminated by postulating language/theory or meaning/belief distinctions of the Carnapian kind. As Quine famously concludes an early essay on Carnap:

“The lore of our fathers is a fabric of sentences...It is a pale grey lore, black with fact and white with convention. But I have found no substantial reasons for concluding that there are any quite black threads in it, or any white ones.”²⁴

In fact, though he did not know it, in making this pragmatist point against Carnap, Quine was recapitulating one of the important ways in which Hegel moves beyond Kant. For Kant, all our empirical activity, cognitive and practical, is discursive activity. In endorsing judgeable contents and practical maxims, knowers and agents are applying concepts. Though further concepts may be developed thereby, for instance by judgments of reflection, one must always

24 “Carnap on Logical Truth”, p. 406 [ref.]

already have concepts in order to be apperceptively *aware* of anything at all. Hegel thought Kant was uncharacteristically, but culpably, uncritical about the origins of our primordial concepts. The locus of those concepts, Hegel thought, lies in *language*, not in some kind of experience understood as prelinguistic. Language, he said, is the existence [Dasein] of Geist—that is, of the whole normatively articulated discursive realm.²⁵ Compare Dewey:

Language in its widest sense—that is, including all means of communication such as, for example, monuments, rituals, and formalized arts—is the medium in which culture exists and through which it is transmitted.²⁶

For Hegel, no less than for Quine and Dewey, we must understand linguistic practices as *both* instituting conceptual norms *and* applying them.²⁷ It is precisely by applying concepts in judging and acting that conceptual content is both made more determinate, going forward, and shows up as always already determinate (in the only sense in which conceptual contents are determinate), looking back.²⁸

²⁵ *Phenomenology of Spirit* [652], [666].

²⁶ *Logic, the Theory of Inquiry, Later Works* Vol. 12, p. 28

²⁷ Here are some characteristic passages:

It is therefore through culture that the individual acquires standing and actuality. His true *original nature* and substance is the alienation of himself as Spirit from his *natural* being. This individuality *moulds* itself by culture into what it intrinsically is. [I: 489]

What, in relation to the single individual, appears as his culture, is the essential moment of the substance itself, viz. the immediate passage of the [mere] thought-form of its universality into actuality; or, culture is the simple soul of the substance by means of which, what is implicit in the substance, acquires an acknowledged, real existence. The process in which the individuality moulds itself by culture is, therefore, at the same time the development of it as the universal, objective essence, i.e. the development of the actual world. Although this world has come into being through individuality, it is for self-consciousness immediately an alienated world which has the form of a fixed and solid reality over against it. [PG 490]

²⁸ See footnote 16.

VI. Rationalism and Pragmatism

Pragmatists who have made the linguistic turn take it that the most important feature of the natural history of creatures like us is that we have *come into language*²⁹: come to engage in distinctively linguistic practices and to exercise distinctively linguistic abilities. This is both an ontogenetic and a phylogenetic achievement. Understanding it requires, at a minimum, addressing three large, interconnected kinds of question. These concern the issues of *demarcation, emergence, and leverage*. The demarcation question is definitional. How are *linguistic* practices and abilities (and hence, the lingualist about discursivity claims, *discursive* ones) to be distinguished from nonlinguistic ones? The emergence question concerns the requirement that any account of language that aspires to being naturalistic in even a very broad sense must explain the possibility of the transition from nonlinguistic to linguistic practices and abilities. How are the abilities we can see in non- or prelinguistic creatures recruited, deployed, and transformed so as to amount to linguistic ones? The leverage question is how to characterize and explain the massive qualitative difference in capacity between linguistic and nonlinguistic creatures: the bonanza of new abilities and possibilities that language opens up for those that do make the transition.

²⁹ We have come to see that there are substantial, potentially controversial presuppositions involved in characterizing this in terms of language *learning*.

One of the principal accomplishments of the classical American pragmatists is the attention they gave to the problem of emergence, to displaying the continuities that make it naturalistically intelligible that species and individuals should be able to cross the boundary separating the prelinguistic from the linguistic. In *Experience and Nature*, Dewey sets the emergence problem this way:

Upon the whole, professed transcendentalists have been more aware than have professed empiricists of the fact that language makes the difference between brute and man. The trouble is that they have lacked naturalistic conception of its origin and status.

30

In his *Logic*, he expands on this thought:

Any theory that rests upon a naturalistic postulate must face the problem of the extraordinary differences that mark off the activities and achievements of human beings from those of other biological forms. It is these differences that have led to the idea that man is completely separated from other animals by properties that come from a non-natural source....The development of language (in its widest sense) out of prior biological activities is, in its connection with wider cultural forces, the key to this transformation. The problem, so viewed, is not the problem of the transition of organic behavior into something wholly discontinuous with it—as is the case when, for example, Reason, Intuition and the A priori are appealed to for explanation of the difference. It is a special form of the general problem of continuity of change and the emergence of new modes of activity—the problem of development at any level.³¹

³⁰ *Experience and Nature, Later Works* Vol. 1, p. 134.

³¹ *Logic, the Theory of Inquiry, Later Works* Vol. 12, p. 50. This emphasis on continuity does not lead Dewey to ignore the differences that language makes:

The hallmark of an untenable intellectualism, he thinks, is an appeal to an inexplicable saltation: the ultimately miraculous dawning of consciousness or self-consciousness, the infusion of reason into a brute. The desire to provide a more satisfactory response to the emergence question than that sort of cartesian approach can offer binds Dewey together with the later Wittgenstein in a common enterprise. The point of many of the toy Sprachspiele the latter describes is to show us how features of discourse that might seem mysterious in a sense that calls for the invocation of a cartesian discontinuity can be exhibited already in practices we can see that intelligent nonlinguistic hominids could master.

When we turn to the demarcation question, however, I think the pragmatists disappoint. What is distinctive of linguistic (or discursive) practices? What sets them apart from prelinguistic or nondiscursive practices? It is one's answer to this question that ties together the emergence question with the leverage question. For the criteria of adequacy for answers to those questions turn on its being the *same* kind of practices and abilities that one has told a story about the nonmiraculous emergence of, in answering the first question, that one then must show can intelligibly account for the huge differences in capabilities, cognitive and practical, that come with the advent of language, in answering the second question. We need not assume that the emergence of language is an all-or-none thing. One might, with Wittgenstein, want to deny that there is or need be a bright line separating the discursive from the nondiscursive, in favor of a

The evidence usually adduced in support of the proposition that lower animals, animals without language, think, turns out, when examined, to be evidence that when men, organisms with power of social discourse, think, they do so with the organs of adaptation used by lower animals, and thus largely repeat in imagination schemes of overt animal action. But to argue from this fact to the conclusion that animals think is like concluding that because every tool, say a plow, originated from some pre-existing natural production, say a crooked root or forked branch, the latter was inherently and antecedently engaged in plowing. The connection is there, but it is the other way around. *Experience and Nature, Later Works* Vol. 1, p. 215.

family-resemblances sort of view. A pluralist-incrementalist response to the demarcation question makes the emergence question easier to answer, but makes the leverage question correspondingly more difficult. I don't think Dewey's metainstrumentalist "tool of tools" line can be made to work to bring the emergence and leverage issues into harmony—but I've argued that elsewhere and won't rehearse my complaints here.³² Apart from that, he seems to offer only vague remarks about language as a making enhanced the possibilities of co-operation and rising above the individual standpoint.³³

I cannot here address the all-important leverage question.³⁴ But the demarcation question is prior. After all, if one is going to say how Geist precipitates out of nature, and how it transforms sentient organisms into sapient ones, one should try to say what it is. The challenge is to offer satisfactory responses to *both* the emergence question *and* the leverage question. Focusing on just one of them makes it too easy. In the passage above, Dewey says in effect that the neo-cartesian intellectualists make the leverage question too easy to respond to, by ignoring (or making it impossible to address) the question of emergence. I have just accused him of making the complementary mistake. In any case, it is clear that the hinge that connects the issues of emergence and leverage is the question of demarcation. For the challenge is to show that the

³² *Experience and Nature, Later Works*, Vol. 1, p. 134. I discuss this approach in Chapters One and Two.

³³ I have in mind passages such as this one:

The importance of language as the necessary, and, in the end, sufficient condition of the existence and transmission of non- purely organic activities and their consequences lies in the fact that, on one side, it is a strictly biological mode of behavior, emerging in natural continuity from earlier organic activities, while, on the other hand, it compels one individual to take the standpoint of other individuals and to see and inquire from a standpoint that is not strictly personal but is common to them as participants or "parties" in a conjoint undertaking. It may be directed by and towards some physical existence. But it first has reference to some other person or persons with whom it institutes communication--the making of something common. Hence, to that extent its reference becomes general and "objective."³³ *Logic, the Theory of Inquiry, Later Works* Vol. 12, p. 52.

³⁴ I do address it in *Making It Explicit* and *Between Saying and Doing*.

same phenomenon that one has accounted for the emergence of can leverage sentience into sapience. So demarcating the realm of linguistic or discursive practices and abilities is an absolutely essential element of the philosophical project I have been describing.

I want to close with a suggestion as to one way fundamental pragmatists (those committed to understanding discursive intentionality as a kind of practical intentionality) who are weak lingualists about discursiveness (take engaging in linguistic practices as a necessary condition of deploying concepts)—a class I take to include at least Peirce, Dewey, the early Heidegger, and the later Wittgenstein—might answer the demarcation question, and so determine definite criteria of adequacy for responses to the leverage question. My idea is that pragmatism can usefully be combined with a *rationalist* criterion of demarcation of the linguistic—and hence of discursiveness in general. By this I mean that what distinguishes the linguistic practice in virtue of which we are sapient and not merely sentient beings is its core practices of giving and asking for *reasons*. A necessary and sufficient condition of being a discursive practice is that some performances are accorded by it the pragmatic significance of *claimings* or assertings. Semantically, claimable or assertible contents are *propositional* contents. Syntactically, what expresses those contents is declarative sentences. This combination of pragmatic, semantic, and syntactic features is the iron triangle of discursiveness. The pragmatist order of explanation of course starts with the pragmatics. The thought is that to have the pragmatic significance of an assertion is to be able both to serve as a reason, and potentially to stand in need of reasons. So propositional contents are those that can play the role both of premise and of conclusion in *inferences*. Discursive practice is accordingly understood as essentially inferentially articulated.

In *Logic: The Theory of Inquiry*, the normative status with which Dewey's pragmatics begins, in terms of which the semantics is to be articulated, is *assertibility*. I have argued on the one hand that to be recognizable as engaging in a practice of making claims and (so) giving and asking for reasons, a community must distinguish at least *two* normative statuses: commitment and entitlement to commitments, and further, that splitting the single status of assertibility into these two aspects pays huge benefits semantically.³⁵ Specifically, one can use them to define three kinds of material inference: commitment-preserving inferences, entitlement-preserving inferences, and incompatibility entailments. The core of my strong inferentialist version of rationalistic pragmatism lies in the claim that conceptual content consists in inferential role in a broad sense, articulated along those three dimensions.³⁶ Of course the underlying rationalist criterion of demarcation of the discursive could be worked out in other ways.

Commitment to a rationalist criterion of demarcation of the discursive requires disagreeing with Wittgenstein: Language *does* have a downtown, and it is the practice of making claims and giving and asking for reasons. Other things we can do with language are ancillary to and parasitic upon these essential core functions. On this view, most of the toy practices Wittgenstein calls "Sprachspiele" are *vocal*, but not genuinely *verbal*, not really *language* games. The builder's utterances in the opening 'Slab' practice, for instance, should not be understood as imperatives. They are vocalizations that have the pragmatic significance of making certain

³⁵ See Chapter Six of my *Articulating Reasons* [Harvard University Press, 2001]

³⁶ For the distinction between weak, strong, and hyperinferentialism, see the Introduction to *Articulating Reasons*. Inferentialism is just one form that rationalism might take. For there is more to reason than inference. Making distinctions, formulating definitions, and producing constructions are all rational processes, alongside drawing conclusions.

responses on the part of the assistant appropriate. But genuine imperatives do that by *saying* what it is that ought to be done. In this full-blooded sense, no practice can contain the genuine imperative “Bring me a slab,” unless it also contains declaratives such as “This is a slab.”

Wittgenstein and Dewey are together in rejecting rationalist criteria of demarcation of the linguistic (and hence the discursive)—indeed, in resisting offering any answer at all to the demarcation question. In Dewey’s case, the idea of a rationalist pragmatism would probably have struck him as a *contradictio in adjecto*. But rationalism as I have described it is not a form of the intellectualism that stands opposed to fundamental pragmatism. It is wholly compatible with understanding discursive intentionality as a kind of practical intentionality: specifically, as the kind that includes practices of making claims and giving and asking for reasons. It aims to say what structure a norm-instituting social practice must have in order properly to be understood as such a practice. It offers a specific proposal for how to understand the kind of practical knowing *how* that adds up to cognitive claiming *that*: it is practical mastery of inferential relations and transitions. And answering the demarcation question about discursive practice in a rationalist manner neither makes it impossible in principle to answer the emergence question nor obliges one to give a cartesian answer to it.

I began my story about pragmatism in an unconventional place: with Kant’s *normative* criterion of demarcation of the discursive, that is, with his idea that what is distinctive of judgments and intentional actions is that they are things we are *responsible* for. They are kinds of *commitments*. But that normative criterion of demarcation was also a *rationalist* criterion of demarcation. For he understood that responsibility, that commitment, as a *rational*

responsibility, as the *justificatory* responsibility to have reasons for ones theoretical and practical commitments, the *ampliative* responsibility to acknowledge their inferential consequences, and the *critical* responsibility to revise commitments that are incompatible, that is, that serve as reasons against one another. Kant's pragmatism consists in his strategy of understanding semantic content in terms of what apperceiving subjects must *do* to fulfill those responsibilities. Judgeable contents have to stand to one another in relations of material consequence and incompatibility: the inferential relations that constrain the process of synthesizing a constellation of commitments and entitlements exhibiting the distinctive unity of apperception. Wittgenstein's example teaches that we should follow Hegel's steps toward naturalizing Kant's notion of norms by understanding norms as implicit in social practices. Normative statuses of responsibility and commitment are *social* statuses: creatures of our practical attitudes of taking or treating each other *as* responsible and committed.

The move beyond Dewey and Wittgenstein to a rationalist, more specifically inferentialist pragmatism that I am recommending is accordingly also a return to pragmatism's roots in German idealism. As Kant synthesized empiricism and rationalism, and the pragmatists synthesized naturalism and empiricism, I'm suggesting that a way forward is to synthesize pragmatism and rationalism—in the form of the rationalist response to the demarcation question.

End

Panoramas of Mind and Meaning

Chapter Three:

Reason, Expression, and the Philosophic Enterprise³⁷

1. In this chapter, I want to address the question: "What is philosophy?"

We might to begin with acknowledge a distinction between things that have *natures* and things that have *histories*. Physical things such as electrons and aromatic compounds would be paradigmatic of the first class, while cultural formations such as English Romantic poetry and Ponzi schemes would be paradigmatic of the second. Applied to the case at hand, this distinction would surely place philosophy on the side of things that have histories. But now we might ask: Does philosophy differ in this respect from physics, chemistry, or biology? Physical, chemical, and biological *things* have natures rather than histories, but what about the disciplines that define and study them? Should physics itself be thought of as something that has a nature, or as something that has a history? Concluding the latter is giving a certain kind of pride of place to the historical. For it is in effect treating the *distinction* between things that have natures and things that have histories, between things studied by the Naturwissenschaften and things studied by the Geisteswissenschaften, as itself a cultural formation: the sort of thing that itself has a history rather than a nature. And from here it is a short step (though not, to be sure, an

³⁷ This chapter was originally published in *What Is Philosophy?*, C.P. Ragland and Sarah Heidt (eds.), Yale University Press, 2001, pp. 74-95.

obligatory one) to the thought that natures themselves are the sort of thing that have a history; certainly the *concepts* electron and aromatic compound are that sort of thing. At this point the door is opened to a thorough-going historicism. It is often thought that this is the point to which Hegel—one of my particular heroes—brought us. I think that thought is correct, as far as it goes, but that we go very wrong if we think that that is where Hegel left us.

To say that philosophy is, at least to begin with, to be understood as the sort of thing that has a history rather than a nature is to foreground the way in which what deserves to be counted as distinctively philosophical activity answers to what has actually been done by those we recognize as precedential, tradition-transforming philosophers. One of Hegel's deepest and most important insights, I think, is indeed that the determinate contentfulness of any universal—in this case, the concept of philosophy—can only be understood in terms of the process by which it incorporates the contingencies of the particulars to which it has actually been applied. But he goes on from there to insist that it is in each case the responsibility of those of us who are heirs to such a conceptual tradition to see to it that it is a *rational* tradition—that the distinction it embodies and enforces between correct and incorrect applications of a concept can be *justified*, that applying it in one case and withholding application in another is something for which *reasons* can be given. It is only insofar as we can do that that we are entitled to understand what we are doing as applying *concepts*. We fulfill that obligation by rationally reconstructing the tradition, finding a coherent, cumulative trajectory through it that reveals it as expressively progressive—as the gradual unfolding into greater explicitness of commitments that can be seen retrospectively as always already having been implicit in it. That is, it is our job to rewrite the history so as to discover in it the revelation of what then retrospectively appears as an antecedent

nature. Hegel balances the insight that even natures have histories by seeing rationality itself as imposing the obligation to construe histories as revelatory of natures.

The aim is to pick out a sequence of precedential instances or applications of a concept that amount to the delineation of a content for the concept, much as a judge at common law is obliged to do. *Making* the tradition rational, is not independent of the labor of concretely *taking* it to be so. It is a criterion of adequacy of each such Whiggish rewriting of our disciplinary history that it create and display continuity and progress by its systematic inclusions and exclusions. The discontinuities that correspond to shifts of topic, the forgetting of lessons, and the degeneration of research programs are invisible from within each such telling; but those differences live on in the spaces between the tellings. Each generation redefines its subject by offering a new retrospective reading of its characteristic concerns and hard-won lessons.³⁸ But also, at any one time there will be diverse interpretations, complete with rival canons, competing designations of heroes, and accounts of their heroic feats. Making canons and baking traditions out of the rich ingredients bequeathed us by our discursive predecessors is a game that all can play.

In this chapter, I am going to sketch one such perspective on what philosophers do—discern a nature as revealed by the history.

³⁸ I am describing, of course, for the concept philosophy an exercise of the sort of recollective rationality (Hegel's "Erinnerung") considered for ordinary determinate empirical concepts in Chapter Three.

Ours is a broadly cognitive enterprise—I say ‘*broadly cognitive*’ to indicate that I mean that philosophers aim at a kind of *understanding*, not, more narrowly, at a kind of *knowledge*. To specify the distinctive sort of understanding that is the characteristic goal of philosophers’ writing is to say what distinguishes that enterprise from that of other sorts of constructive seekers of understanding, such as novelists and scientific theorists. I want to do so by focusing not on the peculiar genre of nonfiction creative writing by which philosophical understanding is typically conveyed (though I think that subject is worthy of consideration), but rather on what is distinctive about the understanding itself: both its particular topic, and its characteristic goal. Philosophy is a self-reflexive enterprise: understanding is not only the *goal* of philosophical inquiry, but its *topic* as well. *We* are its topic; but it is us specifically as *understanding* creatures: *discursive* beings, makers and takers of *reasons*, seekers and speakers of *truth*. Seeing philosophy as addressing the nature and conditions of our rationality is, of course, a very traditional outlook—so traditional, indeed, that it is liable to seem quaint and oldfashioned. I’ll address this issue later, remarking now only that rationalism is one thing, and intellectualism another: pragmatists, too, are concerned with the practices of giving and asking for reasons.

I understand the task of philosophers to have as a central element the explication of concepts—or, put slightly more carefully, the development and application of expressive tools with which to make explicit what is implicit in the use of concepts. When I say "explication of concepts", it is hard not to hear "analysis of meanings." There are obviously affinities between my specification and that which defined the concern specifically of "analytic philosophy" in the

middle years of this century. Indeed, I intend, *inter alia*, to be saying what was right about that conception. But what I have in mind is different in various ways. *Explication*, making explicit, is not the same as *analysis*, at least as that notion was classically conceived. As I use the term, for instance, we have no more privileged access to the contents of our concepts than we do to the facts we use them to state; the concepts and the facts are two sides of one coin.

But the most important difference is that where analysis of meanings is a fundamentally *conservative* enterprise (consider the paradox of analysis), I see the point of explicating concepts rather to be opening them up to rational *criticism*. The rational enterprise, the practice of giving and asking for reasons that lies at the heart of discursive activity, requires not only criticizing *beliefs*, as false or unwarranted, but also criticizing *concepts*. Defective concepts distort our thought and constrain us by limiting the propositions and plans we can entertain as candidates for endorsement in belief and intention. This constraint operates behind our backs, out of our sight, since it limits what we are so much as capable of being aware of. Philosophy, in developing and applying tools for the rational criticism of concepts, seeks to free us from these fetters, by bringing the distorting influences out into the light of conscious day, exposing the commitments implicit in our concepts as vulnerable to rational challenge and debate.

2. The first thing to understand about concepts is that concept is a *normative* concept. This is a lesson we owe ultimately to Kant—the great, gray mother of us all. Kant saw us above all as traffickers in concepts. In fact, in a strict sense, *all* that kantian rational creatures can do is to apply concepts. For that is the genus he took to comprise both *judgment* and *action*, our

theoretical activity and our practical activity. One of Kant's great innovations was his view that what in the first instance distinguishes judgments and actions from the mere behavior of denizens of the realm of nature is that they are things that we are in a distinctive sense *responsible* for. They express *commitments* of ours. The norms or rules that determine what we have committed ourselves to, what we have made ourselves responsible for, by making a judgment or performing an action, Kant calls 'concepts'. Judging and acting involves undertaking commitments whose credentials are always potentially at issue. That is, the commitments embodied in judgments and actions are ones we may or may not be *entitled* to, so that the question of whether they are *correct*, whether they are commitments we *ought* to acknowledge and embrace, can always be raised. One of the forms taken by the responsibility we undertake in judging and acting is the responsibility to give reasons that justify the judgment or the action. And the rules that are the concepts we apply in judging and acting determine what would count as a reason for the judgment and the action.

Commitment, entitlement, responsibility—these are all normative notions. Kant replaces the *ontological* distinction between the physical and the mental with the *deontological* distinction between the realm of nature and the realm of freedom: the distinction between things that merely act regularly and things that are subject to distinctively normative sorts of assessment.

Thus for Kant the great philosophical questions are questions about the source and nature of normativity—of the bindingness or validity [Gültigkeit] of conceptual rules. Descartes had bequeathed to his successors a concern for *certainty*: a matter of our grip on concepts and

ideas—paradigmatically, whether we have a hold on them that is clear and distinct. Kant bequeaths to his successors a concern rather for *necessity*: a matter of the grip concepts have on us, the way they bind or oblige us. ‘Necessary’ [notwendig] for Kant just means “according to a rule”. (That is why he is willing to speak of moral and natural necessity as species of a genus.) The important lesson he takes Hume to have taught isn’t about the threat of skepticism, but about how empirical knowledge is unintelligible if we insist on merely *describing* how things in fact *are*, without moving beyond that to *prescribing* how they *must* be, according to causal rules, and how empirical motivation (and so agency) is unintelligible if we stay at the level of ‘*is*’ and eschew reference to the ‘*ought*’s that outrun what merely is. Looking farther back, Kant finds “the celebrated Mr. Locke” sidetracked into a mere “physiology of the understanding”—the tracing of causal antecedents of thought in place of its justificatory antecedents—through a failure to appreciate the essentially normative character of claims to knowledge. But Kant takes the whole Enlightenment to be animated by an at least implicit appreciation of this point. For mankind’s coming into its intellectual and spiritual majority and maturity consists precisely in taking the sort of personal responsibility for its commitments, both doxastic and practical, insisted upon already by Descartes’ meditator.

This placing of normativity at the center of philosophical concern is the reason behind another of Kant’s signal innovations: the pride of place he accords to *judgment*. In a sharp break with tradition, he takes it that the smallest unit of experience, and hence of awareness, is the judgment. This is because judgments, applications of concepts, are the smallest unit for which knowers can be *responsible*. Concepts by themselves don’t express commitments; they only determine what commitments would be undertaken if they were applied. (Frege will express this

kantian point by saying that judgeable contents are the smallest unit to which pragmatic force—paradigmatically the assertional force that consists in the assertor undertaking a special kind of commitment—can attach. Wittgenstein will distinguish sentences from terms and predicates as the smallest expressions whose free-standing utterance can be used to make a move in a language game.) The most general features of Kant’s understanding of the form of judgment also derive from its role as a unit of responsibility. The “I think” that can accompany all representations (hence being, in its formality, the emptiest of all) is the formal shadow of the transcendental unity of apperception, the locus of responsibility determining a coresponsibility class of concept-applications (including actions), what is responsible *for* its judgments. The objective correlate of this subjective aspect of the form of judgment is the “object=X” to which the judgment is directed, the formal shadow of what the judgment makes the knower responsible *to*.

I think that philosophy is the study of us as creatures who judge and act, that is, as discursive, concept-using creatures. And I think that Kant is right to emphasize that understanding what we do in these terms is attributing to us various kinds of normative status, taking us to be subject to distinctive sorts of normative appraisal. So a central philosophical task is understanding this fundamental normative dimension within which we dwell. Kant’s own approach to this issue, developing themes from Rousseau, is based on the thought that genuinely normative authority (constraint by norms) is distinguished from causal power (constraint by facts) in that it binds only those who *acknowledge* it as binding. Because one is subject only to that authority one subjects oneself to, the normative realm can be understood equally as the realm of *freedom*. So being constrained by norms is not only compatible with freedom—

properly understood, it can be seen to be what freedom consists in. I don't know of a thought that is deeper, more difficult, or more important than this.

3. Kant's most basic idea, I said, is that judgment and action are things we are in a distinctive way *responsible* for. What does it mean to be responsible for them? I think the kind of responsibility in question should be understood to be task responsibility: the responsibility to do something. What (else) do judging and acting oblige us to do? The commitments we undertake by applying concepts in particular circumstances—by judging and acting—are ones we may or may not be entitled to, according to the rules (norms) implicit in those concepts. Showing that we are entitled by the rules to apply the concept in a particular case is *justifying* the commitment we undertake thereby, offering *reasons* for it. That is what we are responsible for, the practical content of our conceptual commitments. In undertaking a conceptual commitment one renders oneself in principle liable to demands for reasons. The normative appraisal to which we subject ourselves in judging and acting is appraisal of our reasons. Further, offering a reason for the application of a concept is always applying another concept: making or rehearsing another judgment or undertaking or acknowledging another practical commitment (Kant's "adopting a maxim"). Conceptual commitments both serve as and stand in need of reasons. The normative realm inhabited by creatures who can judge and act is not only the realm of freedom, it is the realm of reason.³⁹

³⁹ This story is told in more detail in Chapter One.

Understanding the norms for correct application that are implicit in concepts requires understanding the role those concepts play in reasoning: what (applications of concepts) count as reasons for the application of that concept, and what (applications of concepts) the application of that concept counts as a reason for. For apart from such understanding, one cannot fulfill the responsibility one undertakes by making a judgment or performing an action. So what distinguishes concept-using creatures from others is that we know our way around the *space of reasons*. Grasping or understanding a concept just is being able practically to place it in a network of inferential relations: to know what is evidence for or against its being properly applied to a particular case, and what its proper applicability to a particular case counts as evidence for or against. Our capacity to know (or believe) *that* something is the case depends on our having a certain kind of know *how*: the ability to tell what is a reason for what.

The cost of losing sight of this point is to assimilate genuinely conceptual activity, judging and acting, too closely to the behavior of mere animals—creatures who do not live and move and have their being in the normative realm of freedom and reason. We share with other animals (and for that matter, with bits of automatic machinery) the capacity reliably to respond differentially to various kinds of stimuli. We, like they, can be understood as classifying stimuli as being of certain kinds, insofar as we are disposed to produce different repeatable sorts of responses to those stimuli. We can respond differentially to red things by uttering the noise “That is red.” A parrot could be trained to do this, as pigeons are trained to peck at a different button when shown a red figure than when shown a green one. The empiricist tradition is right to emphasize that our capacity to have empirical knowledge begins with and crucially depends on such reliable differential responsive dispositions. But though the story begins with this sort of

classification, it does not end there. For the rationalist tradition is right to emphasize that our classificatory responses count as applications of concepts, and hence as so much as candidates for knowledge, only in virtue of their role in reasoning. The crucial difference between the parrot's utterance of the noise "That is red," and the (let us suppose physically indistinguishable) utterance of a human reporter is that for the latter, but not the former, the utterance has the practical significance of making a claim. Doing that is taking up a normative stance of a kind that can serve as a premise from which to draw conclusions. That is, it can serve as a reason for taking up other stances. And further, it is a stance that itself can stand in need of reasons, at least if challenged by the adoption of other, incompatible stances. Where the parrot is merely responsively sounding off, the human counts as applying a concept just insofar as she is understood as making a move in a game of giving and asking for reasons.

The most basic point of Sellars' rationalist critique of empiricism in his masterwork "Empiricism and the Philosophy of Mind," is that even the *noninferentially* elicited perceptual judgments that the empiricist rightly appreciates as forming the empirical basis for our knowledge can count as judgments (applications of concepts) at all only insofar as they are *inferentially* articulated. Thus the idea that there could be an autonomous language game (a game one could play though one played no other) consisting entirely of noninferentially elicited reports—whether of environing stimuli or of the present contents of one's own mind—is a radical mistake. To apply any concepts *noninferentially*, one must be able also to apply concepts *inferentially*. For it is an essential feature of concepts that their applications can both serve as and stand in need of reasons. Making a report or a perceptual judgment is doing something that essentially, and not just accidentally, has the significance of making available a premise for

reasoning. Learning to observe requires learning to infer. Experience and reasoning are two sides of one coin, two capacities presupposed by concept use that are in principle intelligible only in terms of their relations to each other.⁴⁰

To claim that what distinguishes specifically conceptual classification from classification merely by differential responsive disposition is the inferential articulation of the response—that applications of concepts are essentially what can both serve as and stand in need of reasons—is to assign the game of giving and asking for reasons a preeminent place among discursive practices. For it is to say that what makes a practice *discursive* in the first place is that it incorporates reason-giving practices. Now of course there are many things one can do with concepts besides using them to argue and to justify. And it has seemed perverse to some post-Enlightenment thinkers in any way to privilege the rational, cognitive dimension of language use. But if the tradition I have been sketching is right, the capacity to use concepts in all the other ways explored and exploited by the artists and writers whose imaginative enterprises have rightly been admired by romantic opponents of logocentrism is parasitic on the prosaic inferential practices in virtue of which we are entitled to see concepts as in play in the first place. The game of giving and asking for reasons is not just one game among others one can play with language. It is the game in virtue of the playing of which what one has qualifies as *language* (or thought) at all. I am here disagreeing with Wittgenstein, when he claims that “language has no downtown.” On my view, it does, and that downtown (the region around which all the rest of discourse is arrayed as dependent suburbs, is the practices of giving and asking for reasons. This

⁴⁰ Chapter Seven develops this theme further.

is a kind of linguistic *rationalism*. ‘Rationalism’ in this sense does not entail intellectualism, the doctrine that every *implicit* mastery of a propriety of practice is ultimately to be explained by appeal to a prior *explicit* grasp of a principle. It is entirely compatible with the sort of pragmatism that sees things the other way around.

4. As I am suggesting that we think of them, concepts are broadly inferential norms that implicitly govern practices of giving and asking for reasons. Dummett has suggested a useful model for thinking about the inferential articulation of conceptual contents. Generalizing from the model of meaning Gentzen introduces for sentential operators, Dummett suggests that we think of the use of any expression as involving two components: the circumstances in which it is appropriately used and the appropriate consequences of such use. Since our concern is with the application of the concepts expressed by using linguistic expressions, we can render this as the circumstances of appropriate application of the concept, and the appropriate consequences of such application—that is, what follows from the concept’s being applicable.

Some of the circumstances and consequences of applicability of a concept may be inferential in nature. For instance, one of the circumstances of appropriate application of the concept red is that this concept is applicable wherever the concept scarlet is applicable. And to say that is just another way of saying that the inference from “X is scarlet,” to “X is red,” is a good one. And similarly, one of the consequences of the applicability of the concept red is the applicability of the concept colored. And to say that is just another way of saying that the

inference from “X is red,” to “X is colored,” is a good one. But concepts like red also have *noninferential* circumstances of applicability, such as the visible presence of red things. And concepts such as unjust have noninferential consequences of application—that is, they can make it appropriate to *do* (or not do) something, to make another claim true, not just to *say* or judge that it is true.

Even the immediately empirical concepts of *observables*, which have noninferential *circumstances* of application and the immediately practical *evaluative* concepts, which have noninferential *consequences* of application, however, can be understood to have contents that are inferentially articulated. For all concepts incorporate an implicit commitment to the propriety of the inference from their circumstances to their consequences of application. One cannot use the concept red as including the circumstances and consequences mentioned above without committing oneself to the correctness of the inference from “X is scarlet,” to “X is colored.” So we might decompose the norms that govern the use of concepts into three components: circumstances of appropriate application, appropriate consequences of application, and the propriety of an inference from the circumstances to the consequences. I would prefer to understand the inferential commitment expansively, as including the circumstances and consequences it relates, and so as comprising all three normative elements.

I suggested at the outset that we think of philosophy as charged with producing and deploying tools for the criticism of concepts. The key point here is that concepts may incorporate defective inferences. Dummett offers this suggestive example:

A simple case would be that of a pejorative term, e.g. 'Boche'. The conditions for applying the term to someone is that he is of German nationality; the consequences of its application are that he is barbarous and more prone to cruelty than other Europeans. We should envisage the connections in both directions as sufficiently tight as to be involved in the very meaning of the word: neither could be severed without altering its meaning. Someone who rejects the word does so because he does not want to permit a transition from the grounds for applying the term to the consequences of doing so.⁴¹

(It is useful to focus on a French epithet from the first world war, because we are sufficiently removed from its practical effect to be able to get a theoretical grip on how it works. But the thought should go over *mutatis mutandis* for pejoratives in current circulation.) Dummett's idea is that if you do not accept as correct the inference from German nationality to an unusual disposition to barbarity and cruelty, you can only reject the word. You cannot deny that there are any Boche, for that is just denying that the circumstances of application are ever satisfied, that is, that there are any Germans. And you cannot admit that there are Boche but deny that they are disposed to barbarity and cruelty (this is the "Some of my best friends are Boche," ploy), since that is just taking back in one breath what one has asserted just before. Any use of the term commits the user to the inference that is curled up, implicitly, in it. (At Oscar Wilde's trial the prosecutor read out some passages from the Importance of Being Earnest and said "I put it to you, Mr. Wilde, that this is blasphemy. Is it? Yes or no?" Wilde replied just as he ought on the account I am urging: "Sir, 'blasphemy' is not one of my words."⁴²)

⁴¹ Dummett, Frege: Philosophy of Language [Harper and Row, New York, 1973] p. 454.

⁴² Of course, being right on this point didn't keep Wilde out of trouble, anymore than it did Salman Rushdie.

Although they are perhaps among the most dangerous, it is not just highly-charged words, words that couple ‘descriptive’ circumstances of application with ‘evaluative’ consequences of application that incorporate inferences of which we may need to be critical. The use of *any* expression involves commitment to the propriety of the inference from its circumstances to its consequences of application. These are almost never logically valid inferences. On the contrary, they are what Sellars called “material” inferences: inferences that articulate the content of the concept expressed. Classical disputes about the nature of personal identity, for instance, can be understood as taking the form of arguments about the propriety of such a material inference. We can agree, we may suppose, about the more or less forensic consequences of application of the concept “same person,” having in mind its significance for attributions of (co-)responsibility. When we disagree about the circumstances of application that should be paired with it—for instance whether bodily or neural continuity, or the psychological continuity of memory count for more—we are really disagreeing about the correctness of the inference from the obtaining of these conditions to the ascription of responsibility. The question about what is the correct concept is a question about which inferences to endorse. I think it is helpful to think about a great number of the questions we ask about other important concepts in these same terms: as having the form of queries about what inferences from circumstances to consequences of application we ought to acknowledge as correct, and why. Think in these terms about such very abstract concepts as morally wrong, just, beautiful, true, explain, know, or prove, and again about ‘thicker’ ones such as unkind, cruel, elegant, justify, and understand.

The use of any of these concepts involves a material inferential commitment: commitment to the propriety of a substantial inferential move from the circumstances in which it is appropriate to apply the concept to the consequences of doing so. The concepts are substantive just because the inferences they incorporate are. Exactly this commitment becomes invisible, however, if one conceives conceptual content in terms of *truth conditions*. For the idea of truth conditions is the idea of a single set of conditions that are at once necessary and sufficient for the application of the concept. The idea of individually necessary conditions that are also jointly sufficient is the idea of a set of consequences of application that can also serve as circumstances of application. Thus the circumstances of application are understood as already including the consequences of application, so that no endorsement of a substantive inference is involved in using the concept. The concept of concepts like this is not incoherent. It is the ideal of *logical* or *formal* concepts. Thus it is a criterion of adequacy for introducing logical connectives that they be inferentially conservative: that their introduction and elimination rules be so related that they permit no new inferences involving only the old vocabulary. But it is a bad idea to take this model of the relation between circumstances and consequences of application of logical vocabulary and extend it to encompass also the substantively contentful *nonlogical* concepts that are the currency in which most of our cognitive and practical transactions are conducted.

It is a bad idea because of its built-in conservatism. Understanding meaning or conceptual content in terms of truth conditions—individually necessary and jointly sufficient conditions—squeezes out of the picture the substantive inferential commitment implicit in the use of any nonlogical concept. But it is precisely those inferential commitments that are subject

to *criticism* in the light of substantive collateral beliefs. If one does not believe that Germans are distinctively barbarous or prone to cruelty, then one must not use the concept Boche, just *because* one does not endorse the substantive material inference it incorporates. On the other model, this diagnosis is not available. The most one can say is that one does not know how to specify truth conditions for the concept. But just what is objectionable about it and why does not appear from this theoretical perspective. . Criticism of concepts is always criticism of the inferential connections. For criticizing whether all the individually sufficient conditions (circumstances) “go together”, i.e. are circumstances of application of one concept, just is wondering whether they all have the same consequences of application (and similarly for wondering whether the consequences of application all “go together”).

5. When we think of conceptual contents in the way I am recommending, we can see not only how beliefs can be used to criticize concepts, but also how concepts can be used to criticize beliefs. For it is the material inferences incorporated in our concepts that we use to elaborate the antecedents and consequences of various candidates for belief—to tell what we would be committing ourselves to, what would entitle us to those commitments, what would be incompatible with them, and so on. Once it is accepted that the inferential norms implicit in our concepts are in principle as revisable in the light of evidence as particular beliefs, conceptual and empirical authority appear as two sides of one coin. Rationally justifying our concepts depends on finding out about how things are—about what actually follows from what—as is most evident in the case of massively defective concepts such as Boche.

Adjusting our beliefs in the light of the connections among them dictated by our concepts, and our concepts in the light of our evidence for the substantive beliefs presupposed by the inferences they incorporate, is the rationally reflective enterprise introduced to us by Socrates. It is what results when the rational, normative connections among claims that govern the practice of giving and asking for reasons are themselves brought into the game, as liable to demands for reasons and justification. Saying or thinking something, making it explicit, consists in applying concepts, thereby taking up a stance in the space of reasons, making a move in the game of giving and asking for reasons. The structure of that space, of that game, though, is not given in advance of our finding out how things are with what we are talking about. For what is *really* a reason for what depends on how things *actually* are. But that inferential structure itself can be the subject of claims and thoughts. It can itself be made explicit in the form of claims about what follows from what, what claims are evidence for or against what other claims, what else one would be committing oneself to by making a certain judgment or performing a certain action. So long as the commitment to the propriety of the inference from German nationality to barbarity and unusual cruelty remains merely implicit in the use of term such as ‘Boche’, it is hidden from rational scrutiny. When it is made explicit in the form of the conditional claim “Anyone who is German is barbarous and unusually prone to cruelty,” it is subject to rational challenge and assessment; it can, for instance, be confronted with such counterexamples as Bach and Goethe.

Discursive explicitness, the application of concepts, is Kantian apperception or consciousness. Bringing into discursive explicitness the inferentially articulated conceptual norms in virtue of which we can be conscious or discursively aware of anything at all is the task

of reflection, or self-consciousness. This is the expressive task distinctive of philosophy. Of course, the practitioners of special disciplines, such as membrane physiology, are concerned to unpack and criticize the inferential commitments implicit in using concepts such as lipid soluble with a given set of circumstances and consequences of application, too. It is the emphasis on the “anything at all” distinguishes philosophical reflection from the more focused reflection that goes on within such special disciplines. Earlier I pinned on Kant a view that identifies us as distinctively *rational* creatures, where that is understood as a matter of our being subject to a certain kind of *normative* assessment: we are creatures who can undertake *commitments* and *responsibilities* that are *conceptually* articulated in that their contents are articulated by what would count as *reasons* for them (as well as what other commitments and responsibilities they provide reasons for). One of philosophy’s defining obligations is to supply and deploy an expressive toolbox, filled with concepts that help us make explicit various aspects of *rationality* and *normativity* in general. **The topic of philosophy is normativity in all its guises, and inference in all its forms.** And its task is an *expressive, explicative* one. So it is the job of practitioners of the various philosophical subfields to design and produce specialized expressive tools, and to hone and shape them with use. At the most general level, *inferential* connections are made explicit by *conditionals*, and their *normative* force is made explicit by *deontic* vocabulary. Different branches of philosophy can be distinguished by the different sorts of inference and normativity they address and explicate, the various special senses of “if...then___,” or of ‘ought’ for which they care. Thus philosophers of science, for instance, develop and deploy conditionals codifying causal, functional, teleological, and other explanatory inferential relations, value theorists sharpen our appreciation of the significance of the

differences in the endorsements expressed by prudential, legal, ethical, and aesthetic ‘ought’s, and so on.

6. I said at the outset that I thought of philosophy as defined by its history, rather than by its nature, but that, following Hegel, I think of our task as understanding it by finding or making a nature in or from its history. The gesture I have made in that direction today, though, could be also be summarized in a different kind of definition, namely in the ostensive definition: Philosophy is the kind of thing that Kant and Hegel did (one might immediately want to add Plato, Aristotle, Frege and Wittgenstein to the list, and then we are embarked on the enterprise of turning a gesture into a story, indeed, a history). So one might ask: Why not just say that, and be done with it? While, as I've indicated, I think that specification is a fine place to start, I also think there is a point to trying to be somewhat more explicit about just what sort of thing it is that one takes it Kant and Hegel (and Frege and Wittgenstein) did. Doing that is not being satisfied just with a wave at philosophy as something that has a history. It is trying rationally to reconstruct that tradition, to recast it into a form in which a constellation of ideas can be seen to be emerging, being expressed, refined, and developed.

With those giants, I see philosophy as a discipline whose distinctive concern is with a certain kind of *self-consciousness*: awareness of ourselves as specifically *discursive* (that is, concept-mongering) creatures. It's task is understanding the conditions, nature, and consequences of conceptual norms and the activities—starting with the social practices of giving and asking for reasons—that they make possible and that make them possible. As concept users,

we are beings who can make explicit how things are and what we are doing—even if always only in relief against a background of implicit circumstances, conditions, skills, and practices. Among the things on which we can bring our explicating capacities to bear are those very concept-using capacities that make it possible to make anything at all explicit. Doing that, I am saying, is philosophizing.

It is easy to be misled by the homey familiarity of these sentiments, and correspondingly important to distinguish this characterization from some neighbors with which it is liable to be confused. There is a clear affinity between this view and Kant's coronation of philosophy as "queen of the sciences." For on this account philosophy does extend its view to encompass all activity that is discursive in a broad sense—that is, all activity that presupposes a capacity for judgment and agency, sapience in general. But in this sense, philosophy is at most *a* queen of the sciences, not *the* queen. For the magisterial sweep of its purview does not serve to distinguish it from, say, psychology, sociology, history, literary or cultural criticism, or even journalism. What distinguishes it is the *expressive* nature of its concern with discursiveness in general, rather than its inclusive scope. My sketch was aimed at introducing a specific difference pertaining to philosophy, not a unique privilege with respect to such other disciplines.

Again, as I have characterized it, philosophy does not play a *foundational* role with respect to other disciplines. Its claims do not stand prior to those of the special sciences in some order of ultimate justification. Nor does philosophy sit at the other end of the process as final judge over the propriety of judgments and actions—as though the warrant of ordinary theoretical

and practical applications of concepts remained somehow provisional until certified by philosophical investigation. And philosophy as I have described it likewise asserts no methodological privilege or insight that potentially collides with the actual procedures of other disciplines.

Indeed, philosophy's own proper concerns with the nature of normativity in general, and with its conceptual species in particular, so on inference and justification in general, impinge on the other disciplines in a role that equally well deserves the characterization of "handmaiden." For what we do that has been misunderstood as having foundational or methodological significance is provide and apply tools for unpacking the substantive commitments that are implicit in the concepts deployed throughout the culture, including the specialized disciplines of the high culture. Making those norms and inferences explicit in the form of claims exposes them for the first time to reasoned assessment, challenge, and defense, and so to the sort of rational emendation that is the primary process of conceptual evolution. But once the implicit presuppositions and consequences have been brought out into the daylight of explicitness, the process of assessment, emendation, and so evolution is the business of those whose concepts they are—and not something philosophers have any particular authority over or expertise regarding. Put another way, it is the business of philosophers to figure out ways to increase semantic and discursive self-consciousness. What one does with that self-consciousness is not our business *qua* philosophers—though of course, *qua* intellectuals generally, it may well be.

Philosophy's *expressive* enterprise is grounded in its focus on us as a certain kind of thing, an expressing thing: as at once creatures and creators of conceptual norms, producers and consumers of reasons, beings distinguished by being subject to the peculiar normative force of the better reason. Its concern with us as specifically *normative* creatures sets philosophy off from the empirical disciplines, both the natural and the social sciences. It is this normative character that binds together the currents of thought epitomized in Stanley Cavell's characteristically trenchant aphorism that Kant depsychologized epistemology, Frege depsychologized logic, and Wittgenstein depsychologized psychology. We might add that Hegel depsychologized history. The depsychologizing move in question is equally a desociologizing. For it is a refocusing on the *normative bindingness* of the concepts deployed in ground-level empirical knowledge, reasoning, and thought in general. This is a move beyond the narrowly *natural* (in the sense of the describable order of causes), toward what Hegel called the 'spiritual' [geistig], that is, the *normative* order. That its concern is specifically with our *conceptual* normativity sets philosophy off from the other humanistic disciplines, from the literary as well as the plastic arts. Conceptual commitments are distinguished by their inferential articulation, by the way they can serve as reasons for one another, and by the way they stand in need of reasons, their entitlement always potentially being at issue. Now in asserting the centrality and indispensability, indeed, the criterial role, of practices of giving and asking for reasons, I am far from saying that reasoning—or even thinking—is all anyone ought to do. I am saying that philosophers' distinctive concern is with what else those reason-mongering practices make possible, and how they do, on the one hand, and with what it is that makes them possible—what sort of doings count as sayings, how believing or saying that is founded on knowing how—on the other. It is this distinctive constellation of concerns that makes philosophy the party of

reasons, and philosophers the friends of the norms, the ones who bring out into the light of discursive explicitness our capacity to make things discursively explicit.

End

Panoramas of Mind and Meaning

Chapter Four:

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

I. Intentionality

In this chapter 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."⁴³ 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?⁴⁴

We can specify the content of someone's belief by saying, for instance, that she believes *that* Kant's servant was named 'Lampl'. 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.⁴⁵ 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

⁴³ 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].

⁴⁴ John Searle, *Intentionality* [Cambridge University Press, 1983].

⁴⁵ 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.

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 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.⁴⁶ 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

⁴⁶ I discuss these programs in more detail in Chapter Three of *Between Saying and Doing: Towards an Analytic Pragmatism* [Oxford University Press, 2008].

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.

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.⁴⁷ 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

⁴⁷ *Language, Thought, and other Biological Categories* [MIT Press, 1987].

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 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.⁴⁸

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* 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

⁴⁸ 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.

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.⁴⁹

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.⁵⁰

⁴⁹ *Frege's Philosophy of Language* [New York: Harper and Row, 1973], p. 362.

⁵⁰ "Thought and Talk," in *Inquiries Into Truth and Interpretation* [New York: Oxford University Press, 1984], p. 156.

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, 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).⁵¹ The capacity to make propositionally explicit claims and have conceptually contentful thoughts is intelligible only in the context of implicitly normative social linguistic practices.

⁵¹ Ludwig Wittgenstein *Philosophical Investigations*, G.E.M Anscombe [Wiley-Blackwell, 3rd edition, 1991] §199.

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*. 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.⁵²

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.

⁵² I develop this model further in *Making It Explicit: Reasoning, Representing, and Discursive Commitment* [Harvard University Press, 1994]—especially Chapter Three.

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 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 ($p \& q$)."⁵³ 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 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

⁵³ 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.].

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 Ben 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).⁵⁴

⁵⁴ 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].

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 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

Panoramas of Mind and Meaning

Chapter Five:

Semantic Inferentialism and Logical Expressivism

I: Introduction

In this chapter I want to introduce a way of thinking about semantics that is different from more familiar ones, and on that basis also a new way of thinking about logic. In case that seems insufficiently ambitious, I'll introduce these ideas by sketching a different way of thinking about some important episodes in the history of philosophy, in the era that stretches from Descartes to Kant. I'm going to explain and motivate the two ideas indicated in the title by putting together considerations drawn from three different thinkers: Frege, Dummett, and Sellars or, as I think of them: the sage of Jena, the sage of Oxford, and the sage of Pittsburgh. In each case I'll be picking up strands other than those usually emphasized in reading these figures.

II: Representationalism and Inferentialism

Pre-Kantian empiricists and rationalists alike were notoriously disposed to run together causal and conceptual issues, largely through insufficient appreciation of the normative character of the "order and connection of ideas" that matters for concepts. But there is another, perhaps less appreciated, contrast in play during this period, besides that of the causal and the conceptual, the origin and the justification of our ideas. Enlightenment epistemology was always the home for two somewhat uneasily coexisting conceptions of the conceptual. The fundamental concept of the dominant and characteristic understanding of cognitive contentfulness in the period initiated by Descartes is of course *representation*. However there is a minority semantic tradition that takes *inference* rather than representation as its master concept.

Rationalists such as Spinoza and Leibniz accepted the central role of the concept of representation in explaining human cognitive activity. But they were not prepared to accept Descartes' strategy of treating the possession of representational content as an unexplained explainer—just dividing the world into what is by nature a representing and what by nature can only be represented. Each of them developed instead an account of what it is for one thing to represent another, in terms of the inferential significance of the representing. They were explicitly concerned, as Descartes was not, to be able to explain what it is for something to be understood, taken, treated, or employed *as* a representing *by* the subject: what it is for it to be a representing *to* or *for* that subject (to be "tanquam rem", as if of things, as Descartes puts it). Their idea was that the way in which representings point beyond themselves to something represented is to be understood in terms of *inferential* relations among representings. States and acts acquire content by being caught up in inferences, as premises and conclusions.

Thus a big divide within Enlightenment epistemology concerns the relative explanatory priority accorded to the concepts of representation and inference. The British empiricists were more puzzled than Descartes about representational purport: the property of so much as seeming to be *about* something. But they were clear in seeking to derive inferential relations from the contents of representings, rather than the other way around. In this regard they belong to the still-dominant tradition that reads inferential correctnesses off from representational correctnesses, which are assumed to be antecedently intelligible. That is why Hume could take for granted the contents of his individual representings, but worry about how they could possibly underwrite the correctness of inductive inferences. The post-Cartesian rationalists, the claim is, give rise to a tradition based on a complementary semantically reductive order of explanation. (So Kant, picking up the thread from this tradition, will come to see their involvement in counterfactually robust inferences as essential to empirical representations having the contents that they do.) These *inferentialists* seek to define representational properties in terms of inferential ones, which must accordingly be capable of being understood antecedently. They start with a notion of content as determining what is a *reason* for what, and understand truth and representation as features of ideas that are not only manifested in, but actually *consist* in their role in reasoning. I actually think that the division of pre-Kantian philosophers into representationalists and inferentialists cuts according to deeper principles of their thought than does the nearly coextensional division of them into empiricists and rationalists, though it goes far beyond my brief to argue for that thesis here.

III: Inferentialism and Noninferential Reports

The concepts for which inferential notions of content are least obviously appropriate are those associated with observable properties, such as colors. For the characteristic use of such concepts is precisely in making *noninferential* reports, such as "This ball is red." One of the most important lessons we can learn from Sellars' masterwork, "Empiricism and the Philosophy of Mind" (as from the *Sense Certainty* section of Hegel's *Phenomenology*) is the inferentialist one that even such noninferential reports must be inferentially articulated. Without that requirement, we can't tell the difference between noninferential reporters and automatic machinery such as thermostats and photocells, which also have reliable dispositions to respond differentially to stimuli. What is the important difference between a thermostat that turns the furnace on when the temperature drops to 60 degrees, or a parrot trained to say "That's red," in the presence of red things, on the one hand, and a genuine noninferential reporter of those circumstances, on the other? Each classifies particular stimuli as being of a general kind, the kind, namely, that elicits a repeatable response of a certain sort. In the same sense, of course, a chunk of iron classifies its environment as being of one of two kinds, depending on whether it responds by rusting or not. It is easy, but uninformative, to say that what distinguishes reporters from reliable responders is awareness. In this use, the term is tied to the notion of understanding--the thermostat and the parrot don't understand their responses, those responses mean nothing to them, though they can mean something to us. We can add that the distinction wanted is that between merely responsive classification and specifically *conceptual* classification. The reporter must, as the parrot and thermostat do not, have the *concept* of temperature or cold. It is classifying under such a concept, something the reporter understands or grasps the meaning of, that makes the relevant difference.

It is at this point that Sellars introduces his central thought: that for a response to have *conceptual* content is just for it to play a role in the *inferential* game of making claims and giving and asking for reasons. To grasp or understand such a concept is to have practical mastery over the inferences it is involved in--to know, in the practical sense of being able to distinguish (a kind of know-*how*), what follows from the applicability of a concept, and what it follows from. The parrot doesn't treat "That's red" as incompatible with "That's green", nor as following from "That's scarlet" and entailing "That's colored." Insofar as the repeatable response is not, for the parrot, caught up in practical proprieties of inference and justification, and so of the making of further judgements, it is not a *conceptual* or a *cognitive* matter at all.

It follows immediately from such an inferential demarcation of the conceptual that in order to master *any* concepts, one must master *many* concepts. For grasp of one concept consists in mastery of at least some of its inferential relations to other concepts. Cognitively, grasp of just one concept is the sound of one hand clapping. Another consequence is that to be able to apply one concept *noninferentially*, one must be able to use others *inferentially*. For unless applying it can serve at least as a premise from which to draw inferential consequences, it is not functioning as a concept at all. So the idea that there could be an autonomous language game, one that could be played though one played no other, consisting entirely of noninferential reports (in the case Sellars is most concerned with in *EPM*, even of the current contents of one's own mind) is a radical mistake. (Of course this is compatible with there being languages without theoretical concepts, that is, concepts whose *only* use is inferential. The requirement is that for *any* concepts to have reporting uses, some concepts must have *nonreporting* uses.)

IV: Frege on Begriffliche Inhalt

My purpose at the moment, however, is not to pursue the *consequences* of the inferential understanding of conceptual contents that Sellars recommends, but its *antecedents*. The predecessor it is most interesting to consider is the young Frege. Frege may seem an unlikely heir to this inferentialist tradition. After all, he is usually thought of as the father of the contemporary way of working out the *representationalist* order of explanation, which starts with an independent notion of relations of reference or denotation obtaining between mental or linguistic items and objects and sets of objects in the largely nonmental, nonlinguistic environment, and determines from these in the familiar fashion, first truth conditions for the sentential representations built out of the subsentential ones, and then, from these, a notion of goodness of inference understood in terms of set-theoretic inclusions among the associated sets of truth conditions. But insofar as it is appropriate to read this twentieth century story back into Frege at all, and I am not sure that it is, it would be possible only beginning with the Frege of the 1890's. He starts his semantic investigations, not with the idea of reference, but with that of inference. His seminal first work, the *Begriffsschrift* of 1879, takes as its aim the explication of "conceptual content" [begriffliche Inhalt]. The qualification "conceptual" is explicitly construed in inferential terms:

2] ...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.⁵⁵

Two claims have the same conceptual content iff they have the same inferential role: a good inference is never turned into a bad one by substituting one for the other. This way of specifying the explanatory target to which semantic theories, including referential ones, are directed is picked up by Frege's student Carnap, who in the *Logical Syntax of Language* defines the content of a sentence as the class of non-valid sentences which are its consequences (i.e. can be inferred from it). Sellars in turn picks up the idea from him, as his references to this definition indicate.

By contrast, the tradition Frege initiated in the 1890's makes truth, rather than inference, primary in the order of explanation. Dummett says of this shift:

3] ...in this respect (and [Dummett implausibly but endearingly hastens to add] in this respect alone) Frege's new approach to logic was retrograde. He characterized logic by saying that, while all sciences have truth as their goal, in logic truth is not merely the goal, but the object of study. The traditional answer to the question what is the subject-matter of logic is, however, that it is, not truth, but inference, or, more properly, the relation of logical consequence. This was the received opinion all through the doldrums of logic, until the subject was revitalized by Frege; and it is, surely, the correct view.⁵⁶

And again:

4] It remains that the representation of logic as concerned with a characteristic of sentences, truth, rather than of transitions from sentences to sentences, had highly deleterious effects both in logic and in philosophy. In philosophy it led to a concentration on logical truth and its

⁵⁵ Frege, *Begriffsschrift* (hereafter *BGS*), section 3.

⁵⁶ Dummett, *Frege's Philosophy of Language* [Harper & Row 1973] (hereafter *FPL*), p. 432.

generalization, analytic truth, as the problematic notions, rather than on the notion of a statement's being a deductive consequence of other statements, and hence to solutions involving a distinction between two supposedly utterly different kinds of truth, analytic truth and contingent truth, which would have appeared preposterous and irrelevant if the central problem had from the start been taken to be that of the character of the relation of deductive consequence.⁵⁷

The important thing to realize is that the young Frege has not yet made this false step. Two further points to keep in mind regarding this passage are: first, shifting from concern with inference to concern with truth is one move, understanding truth in terms of prior primitive reference relations is another. Since the mature Frege treats truth as indefinable and primitive, the extraction of a representationalist commitment even from the texts of the 1890's requires further showing (compare Davidson's truth-without-reference view in our own day). Second, understanding the topic of logic in terms of inference is not the same as seeing it in terms of logical inference, or of "deductive consequence", as Dummett puts it (I'll talk about this below under the heading of "formalism" about inference). The view propounded and attributed to Frege below is different, and from the contemporary vantage-point, more surprising, than that Dummett endorses here.

⁵⁷ Dummett, *FPL*, p. 433. A few comments on this passage: First, the "deleterious effects in logic" Dummett has in mind include taking logics to be individuated by their theorems rather than their consequence relations. Although one can do things either way for classical logic, in more interesting cases logics can have the same theorems but different consequence relations. Second, the contrast with analytic is not obviously contingent—why rule out the possibility of necessity that is not conceptual, but, say, physical? Third, the closing claim seems historically wrong. Kant already distinguished analytic from synthetic judgments, and his concerns did not evidently stem from concern with the subject-matter of logic. I include the passage anyway, since I think the shift in emphasis Dummett is endorsing is a good one, although the reasons he advances need filling in and cleaning up.

V: Material Inference

The kind of inference whose correctnesses determine the conceptual contents of its premises and conclusions may be called, following Sellars, *material* inferences. As examples, consider the inference from "Pittsburgh is to the West of Princeton" to "Princeton is to the East of Pittsburgh", and that from "Lightning is seen now" to "Thunder will be heard soon". It is the contents of the concepts West and East that make the first a good inference, and the contents of the concepts lightning and thunder, as well as the temporal concepts, that make the second appropriate. Endorsing these inferences is part of grasping or mastering those concepts, quite apart from any specifically *logical* competence.

Often, however, *inferential* articulation is identified with *logical* articulation. Material inferences are accordingly treated as a derivative category. The idea is that being rational—being subject to the normative force of the better reason, which so puzzled and fascinated the Greeks—can be understood as a purely logical capacity. In part this tendency was encouraged by merely verbally sloppy formulations of the crucial difference between the inferential force of reasons and the physically efficacious force of causes, which render it as the difference between 'logical' and 'natural' compulsion. Mistakes ensue, however, if the concept logical is employed with these circumstances of application conjoined with consequences of application that restrict the notion of logical force of reasons to formally valid inferences. The substantial commitment that is fundamental to this sort of approach is what Sellars calls

5] ...the received dogma...that the inference which finds its expression in "It is raining, therefore the streets will be wet" is an enthymeme.⁵⁸

According to this line of thought, wherever an inference is endorsed, it is because of belief in a conditional. Thus the instanced inference is understood as implicitly involving the conditional "If it is raining, then the streets will be wet". With that "suppressed" premise supplied, the inference is an instance of the formally valid scheme of conditional detachment. The "dogma" expresses a commitment to an order of explanation that treats all inferences as good or bad solely in virtue of their form, with the contents of the claims they involves mattering only for the truth of the (implicit) premises. According to this way of setting things out, there is no such thing as material inference. This view, which understands "good inference" to mean "formally valid inference", postulating implicit premises as needed, might be called a formalist approach to inference. It trades primitive goodnesses of inference for the truth of conditionals. Doing so is taking the retrograde step that Dummett complains about. (It is also what introduces the problem Lewis Carroll exposes in "Achilles and the Tortoise.") The grasp of logic that is attributed must be an implicit grasp, since it need be manifested only in distinguishing material inferences as good and bad, not in any further capacity to manipulate logical vocabulary or endorse tautologies involving them. But what then is the explanatory payoff from attributing such an implicit logical ability rather than just the capacity to assess proprieties of material inference?

The approach Sellars endorses is best understood by reference to the full list of alternatives he considers:

⁵⁸ Sellars "Inference and Meaning," reprinted in *Pure Pragmatics and Possible Worlds* J. Sicha (ed.) [Ridgeview Publishing Co. 1980] (hereafter, *PPPW*), pp. 261/313.

6] ...we have been led to distinguish the following six conceptions of the status of material rules of inference:

(1) Material rules are as essential to meaning (and hence to language and thought) as formal rules, contributing to the architectural detail of its structure within the flying buttresses of logical form.

(2) While not essential to meaning, material rules of inference have an original authority not derived from formal rules, and play an indispensable role in our thinking on matters of fact.

(3) Same as (2) save that the acknowledgment of material rules of inference is held to be a dispensable feature of thought, at best a matter of convenience.

(4) Material rules of inference have a purely derivative authority, though they are genuinely rules of inference.

(5) The sentences which raise these puzzles about material rules of inference are merely abridged formulations of logically valid inferences. (Clearly the distinction between an inference and the formulation of an inference would have to be explored).

(6) Trains of thought which are said to be governed by "material rules of inference" are actually not inferences at all, but rather activated associations which mimic inference, concealing their intellectual nudity with stolen "therefores".⁵⁹

His own position is that an expression has conceptual content conferred on it by being caught up in, playing a certain role in, material inferences:

7] ...it is the first (or "rationalistic") alternative to which we are committed. According to it, material transformation rules determine the descriptive meaning of the expressions of a language within the framework provided by its logical transformation rules... In traditional language, the "content" of concepts as well as their logical "form" is determined by the rules of the Understanding.⁶⁰

⁵⁹ Sellars, "Inference and Meaning" *PPPW* pp. 265/317.

⁶⁰ Sellars, "Inference and Meaning" *PPPW* pp. 284/336.

Should inferentialist explanations begin with inferences pertaining to propositional *form*, or those pertaining to propositional *content*? One important consideration is that the notion of formally valid inferences is definable in a natural way from that of materially correct ones, while there is no converse route. For given a subset of vocabulary that is privileged or distinguished somehow, an inference can be treated as good in virtue of its form, with respect to that vocabulary, just in case

--it is a materially good inference and

-- it cannot be turned into a materially bad one by substituting non-privileged for non-privileged vocabulary, in its premises and conclusions.

Notice that this substitutional notion of formally good inferences need have nothing special to do with *logic*. If it is *logical* form that is of interest, then one must antecedently be able to distinguish some vocabulary as peculiarly logical. That done, the Fregean semantic strategy of looking for inferential features that are invariant under substitution yields a notion of *logically* valid inferences. But if one picks out *theological* (or aesthetic) vocabulary as privileged, then looking at which substitutions of non-theological (or non-aesthetic) vocabulary for non-theological (non-aesthetic) vocabulary preserve material goodness of inference will pick out inferences good in virtue of their *theological* (or aesthetic) form. According to this way of thinking, the *formal* goodness of inferences derives from and is explained in terms of the *material* goodness of inferences, and so ought not to be appealed to in explaining it. Frege's inferentialist way of specifying the characteristic linguistic role in virtue of which vocabulary qualifies as logical is discussed below.

VI: Elucidative Rationality

So far I have indicated briefly two related claims: that conceptual contents are inferential roles, and that the inferences that matter for such contents in general must be conceived to include those that are in some sense *materially correct*, not just those that are *formally valid*. I'll argue in a moment that a commitment to the second of these, no less than the first, is to be found already in Frege's early writings, though not in the developed form to which Sellars brings it. But in both thinkers these ideas are combined with a third, which I believe makes this line of thought especially attractive. In one of his early papers, Sellars introduces the idea this way:

8] Socratic method serves the purpose of making explicit the rules we have adopted for thought and action, and I shall be interpreting our judgments to the effect that A causally necessitates B as the expression of a rule governing our use of the terms 'A' and 'B'.⁶¹

Sellars understands such modal statements as inference licenses, which formulate as the content of a claim the appropriateness of inferential transitions. More than this, he understands the function of such statements to be making explicit, in the form of assertible rules, commitments that had hitherto remained implicit in inferential practices. Socratic method is a way of bringing our practices under rational control, by expressing them explicitly in a form in which they can be confronted with objections and alternatives, a form in which they can be exhibited as the conclusions of inferences seeking to justify them on the basis of premises advanced as reasons, and as premises in further inferences exploring the consequences of accepting them.

⁶¹ Sellars, "Language, Rules, and Behavior" footnote 2 to p. 136/296 in *PPPW*.

In the passage just quoted, Sellars tells us that the enterprise within which we ought to understand the characteristic function of inference licenses is a form of rationality that centers on the notion of *expression*: making *explicit* in a form that can be thought or said, what is *implicit* in what is done. This is a dark and pregnant claim, but I believe it epitomizes a radical and distinctive insight. In what follows I hope to shed some light on it and its role in an inferentialist vision of things. The general idea is that the paradigmatically rational process that Sellars invokes under the heading of "Socratic method" depends upon the possibility of making implicit commitments explicit in the form of claims. *Expressing* them in this sense is bringing them into the game of giving and asking for reasons as playing the special sort of role in virtue of which something has a conceptual content at all: namely an inferential role, as premise and conclusion of inferences. This sort of rationality is distinct from, but obviously related to the sort of rationality that then consists in making the appropriate inferential moves. Even totalitarian versions of the latter, for instance those that would assimilate all goodness of inference to logical validity, or to instrumental prudence (that is, efficiency at getting what one wants), depend upon the possibility of expressing considerations in a form in which they can be given as reasons, and reasons demanded for them. All the more does Socratic reflection on our practices, particularly on those material-inferential practices that determine the conceptual contents of thoughts and beliefs, depend on the possibility of their explicit expression.

VII: Frege on the Expressive Role of Logic

To begin to explicate this notion of explication, it is helpful to return to the consideration of the young Frege's inferentialist program. Frege's *Begriffsschrift* is remarkable not only for the inferential idiom in which it specifies its topic, but equally for how it conceives its relation to that topic. The task of the work is officially an expressive one; not to prove something, but to say something. Frege's logical notation is designed for expressing conceptual contents, making explicit the inferential involvements that are implicit in anything that possesses such content. As passage [2] quoted above puts it: "Whatever is needed for a correct inference is fully expressed". Talking about this project, Frege says:

9] Right from the start I had in mind the expression of a content...But the content is to be rendered more exactly than is done by verbal language... Speech often only indicates by inessential marks or by imagery what a concept-script should spell out in full.⁶²

The concept-script is a formal language for the explicit codification of conceptual contents. In the Preface to BGS, Frege laments that even in science concepts are formed haphazardly, so that the ones employing them are scarcely aware of what they mean, of what their content really is. When the correctness of particular inferences is at issue, this sort of unclarity may preclude rational settlement of the issue. What is needed is a notation within which the rough-and-ready conceptual contents of the sciences, beginning with mathematics, can be reformulated so as to wear their contents on their sleeves. The explanatory target here avowedly concerns a sort of inference, not a sort of truth, and the sort of inference involved is content-conferring material inferences, not the derivative formal ones.

⁶² Frege, from "Boole's logical Calculus and the Concept-script", *Posthumous Writings* (hereafter *PW*) pp.12-13.

Frege explicitly contrasts his approach with that of those, such as Boole, who conceive their formal language only in terms of formal inference, and so express no material contents:

10] The reason for this inability to form concepts in a scientific manner lies in the lack of one of the two components of which every highly developed language must consist. That is, we may distinguish the formal part...from the material part proper. The signs of arithmetic correspond to the latter. What we still lack is the logical cement that will bind these building stones firmly together...In contrast, Boole's symbolic logic only represents the formal part of the language.⁶³

By contrast:

- 11] 1. My concept-script has a more far-reaching aim than Boolean logic, in that it strives to make it possible to present a content when combined with arithmetical and geometrical signs...
2. Disregarding content, within the domain of pure logic it also, thanks to the notation for generality, commands a somewhat wider domain...
4. It is in a position to represent the formation of the concepts actually needed in science...⁶⁴

It is the wider domain to which his expressive ambition extends that Frege sees as characteristic of his approach. Since contents are determined by inferences, expressing inferences explicitly will permit the expression of any sort of content at all:

12] It seems to me to be easier still to extend the domain of this formula language to include geometry. We would only have to add a few signs for the intuitive relations that occur

⁶³ Frege, *PW* p. 13.

⁶⁴ Frege, *PW* p. 46.

there...The transition to the pure theory of motion and then to mechanics and physics could follow at this point.⁶⁵

Frege's early understanding of logic offers some specific content to the notion of explicitly expressing what is implicit in a conceptual content, which is what is required to fill in a notion of expressive or elucidating rationality that might be laid along side (and perhaps even be discovered to be presupposed by) notions of rationality as accurate representation, as logically valid inference, and as instrumental practical reasoning. Before the fateful step from seeing logic as an attempt to codify inferences to seeing it as the search for a special kind of truth is made, which Dummett bemoans, Frege's aim is to introduce vocabulary that will let one *say* (explicitly) what otherwise one can only *do* (implicitly). Consider the conditional, with which the *Begriffsschrift* begins. Frege says of it:

13] The precisely defined hypothetical relation between contents of possible judgments [Frege's conditional] has a similar significance for the foundation of my concept-script to that which identity of extensions has for Boolean logic.⁶⁶

[I think it is hard to overestimate the importance of this passage in understanding what is distinctive about Frege's *Begriffsschrift* project. After all, contemporary Tarskian model-theoretic semantics depends precisely on relations among extensions. Frege is saying that his distinctive idea—in what is, after all, the founding document of modern formal logic—is to do things otherwise.] Why the conditional? Prior to the introduction of such a conditional locution, one could *do* something, one could treat a judgement as having a certain content (implicitly attribute that content to it) by endorsing various inferences involving it and rejecting others. After conditional locutions have been introduced, one can *say*, as part of the content of a claim (something that can serve as a premise and conclusion in inference), *that* a certain inference is

⁶⁵ Frege, *Begriffsschrift* Preface, in van Heijenoort (ed.) *From Frege to Godel* Harvard Press, 1967 p. 7.

⁶⁶ Frege, *PW* p. 16.

acceptable. One is able to make explicit material inferential relations between an antecedent or premise and a consequent or conclusion. Since according to the inferentialist view of conceptual contents, it is these implicitly recognized material inferential relations that conceptual contents consist in, the conditional permits such contents to be explicitly expressed. If there is a disagreement about the goodness of an inference, it is possible to say what the dispute is about, and offer reasons one way or the other. The conditional is the paradigm of a locution that permits one to make inferential commitments explicit as the contents of judgments. In a similar fashion, introducing negation makes it possible to express explicitly material incompatibilities of sentences, which also contribute to their content. The picture is accordingly one whereby first, formal validity of inferences is defined in terms of materially correct inferences and some privileged vocabulary; second, that privileged vocabulary is identified as logical vocabulary; and third, what it is for something to be a bit of logical vocabulary is explained in terms of its semantically expressive role.

Frege is not as explicit about the role of materially correct inferences as Sellars is, but his commitment to the notion is clear from the relation between two of the views that have been extracted from the *Begriffsschrift*: expressivism about logic and inferentialism about content. Expressivism about logic means that Frege treats logical vocabulary as having a distinctive expressive role--making explicit the inferences that are implicit in the conceptual contents of nonlogical concepts. Inferentialism about those conceptual contents is taking them to be identified and individuated by their inferential roles. Together these views require that it be coherent to talk about inference prior to the introduction of specifically logical vocabulary, and so prior to the identification of any inferences as good in virtue of their form. In the context of an inferential understanding of conceptual contents, an expressivist approach presupposes a notion of nonlogical inference, the inferences in virtue of which concepts have nonlogical content. Thus the young Frege envisages a field of material inferences that confer conceptual

content on sentences caught up in them. So although Frege does not offer an explanation of the concept, in the *Begriffsschrift* his expressive, explicating project commits him to something playing the role Sellars later picks out by the phrase "material inference".

VIII: Dummett's Model, and Gentzen

So far three themes have been introduced:

--that conceptual content is to be understood in terms of role in reasoning rather than exclusively in terms of representation,

--that the capacity for such reasoning is not to be identified exclusively with mastery of a logical calculus, and

--that besides theoretical and practical reasoning using contents constituted by their role in material inferences, there is a kind of expressive rationality that consists in making implicit content-conferring inferential commitments explicit as the contents of assertible commitments. In this way, the material inferential practices, which govern and make possible the game of giving and asking for reasons, are brought into that game, and so into consciousness, as explicit topics of discussion and justification.

These three themes, to be found in the early works of both Frege and Sellars, provide the beginnings of the structure within which modern inferentialism develops. These ideas can be made more definite by considering a general model of conceptual contents as inferential roles that has been recommended by Dummett. According to that model, the use of any linguistic expression or concept has two aspects: the *circumstances* under which it is correctly applied, uttered, or used, and the appropriate *consequences* of its application, utterance, or use. Though Dummett does not make this point, this model can be connected to inferentialism via the principle that the content to which one is committed by using the concept or expression may be represented by the inference one implicitly endorses by such use, the inference, namely, from the circumstances of appropriate employment to the appropriate consequences of such employment.

The original source for the model lies in a treatment of the grammatical category of sentential connectives. Dummett's two-aspect model is a generalization of a standard way of specifying the inferential roles of logical connectives, due ultimately to Gentzen. Gentzen famously defined connectives by specifying introduction rules, or inferentially sufficient conditions for the employment of the connective, and elimination rules, or inferentially necessary consequences of the employment of the connective. So, to define the inferential role of an expression '&' of Boolean conjunction, one specifies that anyone who is committed to p , and committed to q , is thereby to count also as committed to $p \& q$, and that anyone who is committed to $p \& q$ is thereby committed both to p and to q . The first schema specifies, by means of expressions that do not contain the connective, the circumstances under which one is committed to claims expressed by sentences that do contain (as principle connective) the connective whose inferential role is being defined, that is, the sets of premises that entail them. The second schema specifies, by means of expressions that do not contain the connective, the consequences of being committed to claims expressed by sentences that do contain (as principle connective) the connective whose inferential role is being defined, that is, the sets of consequences that they entail.

IX: Circumstances and Consequences for Sentences

Dummett makes a remarkable contribution to inferentialist approaches to conceptual content by showing how this model can be generalized from logical connectives to provide a uniform treatment of the meanings of expressions of other grammatical categories, in particular sentences, predicates and common nouns, and singular terms. The application to the propositional contents expressed by whole sentences is straightforward. What corresponds to an *introduction* rule for a propositional content is the set of *sufficient* conditions for asserting it, and what corresponds to an *elimination* rule is the set of *necessary* consequences of asserting it, that is, what follows from doing so. Dummett says:

14] Learning to use a statement of a given form involves, then, learning two things: the conditions under which one is justified in making the statement; and what constitutes acceptance of it, i.e., the consequences of accepting it.⁶⁷

Dummett presents his model as specifying two fundamental features of the *use* of linguistic expressions, an idea I'll return to below. In what follows here, though, I'll be applying it in the context of the previous ideas to bring into relief the implicit material inferential *content* a concept or expression acquires in virtue of being used in the ways specified by these two 'aspects'. The link between pragmatic significance and inferential content is supplied by the fact that asserting a sentence is implicitly undertaking a commitment to the correctness of the material inference from its circumstances to its consequences of application.

⁶⁷ Dummett, *FPL* p. 453.

Understanding or grasping a propositional content is here presented not as the turning on of a Cartesian light, but as practical mastery of a certain kind of inferentially articulated doing: responding differentially according to the circumstances of proper application of a concept, and distinguishing the proper inferential consequences of such application. This is not an all-or-none affair; the metallurgist understands the concept tellurium better than I do, for training has made her master of the inferential intricacies of its employment in a way that I can only crudely approximate. Thinking clearly is on this inferentialist rendering a matter of knowing what one is committing oneself to by a certain claim, and what would entitle one to that commitment. Writing clearly is providing enough clues for a reader to infer what one intends to be committed to by each claim, and what one takes it would entitle one to that commitment. Failure to grasp either of these components is failure to grasp the inferential commitment use of the concept involves, and so failure to grasp its conceptual content.

Failure to think about both the circumstances and consequences of application leads to semantic theories that are literally one-sided. Verificationists, assertibilists, and reliabilists make the mistake of treating the *first* aspect as exhausting content. Understanding or grasping a content is taken to consist in practically mastering the circumstances under which one becomes entitled or committed to endorse a claim, quite apart from any grasp of what one becomes entitled or committed to by such endorsement. But this cannot be right. For claims can have the same circumstances of application and different consequences of application, as for instance 'I foresee that I will write a book about Hegel' and 'I will write a book about Hegel' do. We can at least regiment a use of 'foresee' that makes the former sentence have just the same assertibility conditions as the latter. But substituting the one for the other turns the very safe conditional "If I will write a book about Hegel, then I will write a book about Hegel," into the risky "If I *foresee* that I will write a book about Hegel, then I will write a book about Hegel." The possibility that I

might be hit by a bus does not affect the assessment of the inference codified by the first conditional, but is quite relevant to the assessment of the second inference.

And the point of the discussion of Sellars' application of inferentialist ideas to the understanding of noninferential reports, at the beginning of this essay, was that parrots and photocells and so on might reliably discriminate the circumstances in which the concept 'red' should be applied, without thereby grasping that concept, precisely in the case where they have no mastery of the consequences of such application—when they can't tell that it follows from something being red that it is colored, that it is not a prime number, and so on. You do not convey to me the content of the concept 'gleeb' by supplying me with an infallible gleebness tester, which lights up when and only when exposed to gleeb things. I would in that case know what things were gleeb, without knowing what I was saying about them when I called them that, what I had found out about them or committed myself to. Dummett offers two examples of philosophically important concepts where it is useful to be reminded of this point:

15] An account, however accurate, of the conditions under which some predicate is rightly applied may thus miss important intuitive features of its meaning; in particular, it may leave out what we take to be the point of our use of the predicate. A philosophical account of the notion of truth can thus not necessarily be attained by a definition of the predicate 'true', even if one is possible, since such a definition may be correct only in the sense that it specifies correctly the application of the predicate, while leaving the connections between this predicate and other notions quite obscure.⁶⁸

Even more clearly:

16] A good example would be the word 'valid' as applied to various forms of argument. We might reckon the syntactic characterization of validity as giving the criterion for applying the

⁶⁸ Dummett, *FPL* p. 455.

predicate 'valid' to an argument, and the semantic characterization of validity of giving the consequences of such an application. ...if he is taught in a very unimaginative way, he may see the classification of arguments into valid and invalid ones as resembling the classification of poems into sonnets and non-sonnets, and so fail to grasp that the fact that an argument is valid provides any grounds for accepting the conclusion if one accepts the premises. We should naturally say that he had missed the point of the distinction.⁶⁹

Pragmatists of the classical sort, on the other hand, make the converse mistake of identifying propositional contents exclusively with the *consequences* of endorsing a claim, looking downstream to the claim's role as a premise in practical reasoning and ignoring its proper antecedents upstream. [For present purposes, that the emphasis is on *practical* consequences doesn't matter.] Yet one can know what follows from the claim that someone is responsible for a particular action, that an action is immoral or sinful, that a remark is true or in bad taste, without for that reason counting as understanding the claims involved, if one has no idea when it is appropriate to make those claims or apply those concepts. Being classified as AWOL does have the consequence that one is liable to be arrested, but the specific circumstances under which one acquires that liability are equally essential to the concept.

⁶⁹ Dummett, *FPL* pp. 453-4.

X: 'Derivation', Prior, Belnap, and Conservativeness

Of course, such one-sided theories don't simply ignore the aspects of content they don't treat as central. Dummett says:

17] ...most philosophical observations about meaning embody a claim to perceive... a simple pattern: the meaning of a sentence consists in the conditions for its truth and falsity, or in the method of its verification, or in the practical consequences of accepting it. Such dicta cannot be taken to be so naive as to involve overlooking the fact that there are many other features of the use of a sentence than the one singled out as being that in which its meaning consists: rather, the hope is that we shall be able to give an account of the connection that exists between the different aspects of meaning. One particular aspect will be taken as central, as constitutive of the meaning of any given sentence...; all other features of the use of the sentence will then be explained by a uniform account of their derivation from that feature taken as central.⁷⁰

I think this is a very helpful way to think about the structure of theories of meaning in general, but two observations should be made. First, the principle that the task of a theory of meaning is to explain the use of expressions to which meanings are attributed does not mandate identifying meaning with an aspect of use. Perhaps meanings are to use as theoretical entities are to the observable ones whose antics they are postulated to explain. We need not follow Dummett in his semantic instrumentalism. Second, one might deny that there are meanings in this sense: that is deny that all the features of the use of an expression can be derived in a uniform way from anything we know about it. Dummett suggests that this is the view of the later Wittgenstein. One who takes language to be a motley in this sense will deny that there are such a things as meanings to be the objects of a theory (without, of course, denying that expressions are meaningful). Keeping these caveats in mind, we will find that pursuing this notion of *derivation* provides a helpful perspective on the idea of conceptual contents articulated according to

⁷⁰ Dummett, *FPL* pp. 456-7.

material inferences, and on the role of explicit inference licenses such as conditional statements in expressing and elucidating such inferences, and so such contents.

For the special case of defining the inferential roles of logical connectives by pairs of sets of rules for their introduction and for their elimination, which motivates Dummett's broader model, there is a special condition it is appropriate to impose on the relation between the two sorts of rules.

18] In the case of a logical constant, we may regard the introduction rules governing it as giving conditions for the assertion of a statement of which it is the main operator, and the elimination rules as giving the consequences of such a statement: the demand for harmony between them is then expressible as the requirement that the addition of the constant to a language produces a conservative extension of that language.⁷¹

Recognition of the appropriateness of such a requirement arises from consideration of connectives with 'inconsistent' contents. As Prior pointed out, if we define a connective, which after Belnap we may call 'tonk', as having the introduction rule proper to disjunction and the elimination rule proper to conjunction, then the first rule licenses the transition from p to p tonk q , for arbitrary q , and the second licenses the transition from p tonk q to q , and we have what he called a "runabout inference ticket" permitting any arbitrary inference. Prior thought that this possibility shows the bankruptcy of Gentzen-style definitions of inferential roles. Belnap shows rather that when logical vocabulary is being introduced, one must constrain such definitions by the condition that the rule not license any inferences involving only old vocabulary that were not already licensed before the logical vocabulary was introduced, that is, that the new rules provide an inferentially conservative extension of the original field of inferences. Such a constraint is necessary and sufficient to keep from getting into trouble with Gentzen-style definitions. But the

⁷¹ Dummett, *FPL* p. 454.

expressive account of what distinguishes logical vocabulary shows us a deep *reason* for this demand; it is needed not only to avoid horrible consequences but because otherwise logical vocabulary cannot perform its expressive function. Unless the introduction and elimination rules are inferentially conservative, the introduction of the new vocabulary licenses new material inferences, and so alters the contents associated with the old vocabulary. So if logical vocabulary is to play its distinctive expressive role of making explicit the original material inferences, and so conceptual contents expressed by the old vocabulary, it must be a criterion of adequacy for introducing logical vocabulary that no new inferences involving only the old vocabulary be made appropriate thereby.

XI: 'Boche' and the Elucidation of Inferential Commitments

The problem of what Dummett calls a lack of 'harmony' between the circumstances and the consequences of application of a concept may arise for concepts with material contents, however. Seeing how it does provides further help in understanding the notion of expressive rationality, and the way in which the explicating role of logical vocabulary contributes to the clarification of concepts. For conceptual change can be:

19] ...motivated by the desire to attain or preserve a harmony between the two aspects of an expression's meaning. A simple case would be that of a pejorative term, e.g. 'Boche'. The conditions for applying the term to someone is that he is of German nationality; the consequences of its application are that he is barbarous and more prone to cruelty than other Europeans. We should envisage the connections in both directions as sufficiently tight as to be involved in the very meaning of the word: neither could be severed without altering its meaning. Someone who rejects the word does so because he does not want to permit a transition from the grounds for applying the term to the consequences of doing so. The addition of the term 'Boche' to a language which did not previously contain it would produce a non-conservative extension, i.e. one in which certain other statements which did not contain the term were inferable from other statements not containing it which were not previously inferable...⁷²

This crucial passage makes a number of points that are worth untangling. First of all, it shows how concepts can be criticized on the basis of substantive beliefs. If one does not believe that the inference from German nationality to cruelty is a good one, one must eschew the concept or expression "Boche". For one cannot deny that there are any Boche--that is just denying that anyone is German, which is patently false. One cannot admit that there are Boche and deny that they are cruel--that is just attempting to take back with one claim what one has committed

⁷² Dummett, *FPL* p. 454.

oneself to with another. One can only refuse to employ the concept, on the grounds that it embodies an inference one does not endorse.

The prosecutor at Oscar Wilde's trial at one point read out some of the more hair-raising passages from "The Importance of Being Earnest" and said "I put it to you, Mr. Wilde, that this is *blasphemy*. Is it or is it not?" Wilde made exactly the reply he ought to make—indeed, the only one he could make—given the considerations being presented here and the circumstances and consequences of application of the concept in question. He said "Sir, 'blasphemy' is not one of my words."

Highly charged words such as "nigger", "whore", "faggot", "lady", "Communist", "Republican", have seemed to some a special case because they couple 'descriptive' circumstances of application to 'evaluative' consequences. But this is not the only sort of expression embodying inferences that requires close scrutiny. The use of any concept or expression involves commitment to an inference from its grounds to its consequences of application. Critical thinkers, or merely fastidious ones, must examine their idioms to be sure that they are prepared to endorse and so defend the appropriateness of the material inferential transitions implicit in the concepts they employ. In Reason's fight against thought debased by prejudice and propaganda, the first rule is that potentially controversial material inferential commitments should be made explicit as claims, exposing them both as vulnerable to reasoned challenge and as in need of reasoned defense. They must not be allowed to remain curled up inside loaded phrases such as "enemy of the people" or "law and order."

It is in this process that formal logical vocabulary such as the conditional plays its explicating role. It permits the formulation, as explicit claims, of the inferential commitments that otherwise remain implicit and unexamined in the contents of material concepts. Logical locutions make it possible to display the relevant grounds and consequences, and to assert their inferential relation. Formulating as an explicit claim the inferential commitment implicit in the content brings it out into the open as liable to challenges and demands for justification, just as with any assertion. In this way explicit expression plays an elucidating role, functioning to groom and improve our inferential commitments, and so our conceptual contents--a role, in short, in the practices of reflective rationality or "Socratic method".

But if Dummett is suggesting that what is wrong with the concept 'Boche' (or 'nigger') is that its addition represents a nonconservative extension of the rest of the language, he is mistaken. Its nonconservativeness just shows that it has a substantive content, in that it implicitly involves a material inference that is not already implicit in the contents of other concepts being employed. Outside of logic, this is no bad thing. Conceptual progress in science often consists in introducing just such novel contents. The concept of temperature was introduced with certain criteria or circumstances of appropriate application, and certain consequences of application. As new ways of measuring temperature are introduced, and new theoretical and practical consequences of temperature measurements adopted, the complex inferential commitment that determines the significance of using the concept of temperature evolves.

The proper question to ask in evaluating the introduction and evolution of a concept is not whether the inference embodied is one that is already endorsed, so that no new content is really involved, but rather whether that inference is one that *ought* to be endorsed. The problem

with 'Boche' or 'nigger' is not that once we explicitly confront the material inferential commitment that gives them their content, it turns out to be novel, but that it can then be seen to be indefensible and inappropriate--a commitment we cannot become entitled to. We want to be aware of the inferential commitments our concepts involve, to be able to make them explicit, and to be able to justify them. But there are other ways of justifying them than showing that we were already implicitly committed to them, before introducing or altering the concept in question.

XII: Harmony and Material Inference

Even in the cases where it does make sense to identify harmony of circumstances and consequences with inferential conservativeness, the attribution of conservativeness is always relative to a background set of material inferential practices, the ones that are conservatively extended by the vocabulary in question. Conservativeness is a property of the conceptual content only in the context of other contents, not something it has by itself. Thus there can be pairs of logical connectives, either of which is all right by itself, but both of which cannot be included in a consistent system. It is a peculiar ideal of harmony that would be realized by a system of conceptual contents such that the material inferences implicit in every subset of concepts represented a conservative extension of the remaining concepts, in that no inferences involving only the remaining ones are licensed that are not licensed already by the contents associated just with those remaining concepts. Such a system is an idealization, because all of its concepts would already be out in the open; none remaining hidden, to be revealed only by drawing conclusions from premises that have never been conjoined before, following out unexplored lines of reasoning, drawing consequences one was not previously aware one would be entitled or committed to by some set of premises. In short, this would be a case where Socratic reflection, making implicit commitments explicit and examining their consequences and possible justifications, would never motivate one to alter contents or commitments. Such complete transparency of commitment and entitlement is in some sense an ideal projected by the sort of Socratic practice that finds current contents and commitments wanting by confronting them with each other, pointing out inferential features of each of which we were unaware. But as Wittgenstein teaches in general, it should not be assumed that our scheme is like this, or depends upon an underlying set of contents like this, just because we are obliged to remove any particular ways in which we discover it to fall short.

These are reasons to part company with the suggestion, forwarded in the passage above, that inferential conservatism is a necessary condition of a 'harmonious' concept--one that won't 'tonk up' a conceptual scheme. In a footnote, Dummett explicitly denies that conservativeness can in general be treated as a sufficient condition of harmony:

20] This is not to say that the character of the harmony demanded is always easy to explain, or that it can always be accounted for in terms of the notion of a conservative extension. ...the most difficult case is probably the vexed problem of personal identity.⁷³

In another place, this remark about personal identity is laid out in more detail:

21] We have reasonably sharp criteria which we apply in ordinary cases for deciding questions of personal identity: and there are also fairly clear consequences attaching to the settlement of such a question one way or the other, namely those relating to ascriptions of responsibility, both moral and legal, to the rights and obligations which a person has... What is much harder is to give an account of the connection between the criteria for the truth of a statement of personal identity and the consequences of accepting it. We can easily imagine people who use different criteria from ours...Precisely what would make the criteria they used criteria for personal identity would lie in their attaching the same consequence, in regard to responsibility, motivation, etc., to their statements of personal identity as we do to ours. If there existed a clear method for deriving, as it were, the consequences of a statement from the criteria for its truth, then the difference between such people and ourselves would have the character of a factual disagreement, and one side would be able to show the other to be wrong. If there were no connection between truth-grounds and consequences, then the disagreement between us would lie merely in a preference for different concepts, and there would be no right or wrong in the matter at all.⁷⁴

⁷³ Dummett, *FPL* p. 455n.

⁷⁴ Dummett, *FPL* p. 358.

Dummett thinks that there is a general problem concerning the way in which the circumstances and consequences of application of expressions or concepts ought to fit together. Some sort of 'harmony' seems to be required between these two aspects of the use. The puzzling thing, he seems to be saying, is that the harmony required cannot happily be assimilated either to compulsion by facts or to the dictates of freely chosen meanings. But the options: matter of fact or relation of ideas, expression of commitment as belief or expression of commitment as meaning are not ones that readers of "Two Dogmas of Empiricism" ought to be tempted to treat as exhaustive.

The notion of a completely factual issue that Dummett appeals to in this passage is one in which the applicability of a concept is settled straightforwardly by the application of other concepts, the concepts that specify the necessary and sufficient conditions that determine the truth conditions of claims involving the original concept. This conception, envisaged by a model of conceptual content as necessary and sufficient conditions, seems to require a conceptual scheme that is ideally transparent in the way mentioned above, in that it is immune to Socratic criticism. For that conception insists that these coincide in that the jointly sufficient conditions already entail the individually necessary ones, so that it is attractive to talk about content as truth conditions, rather than focussing on the substantive inferential commitments that relate the sufficient to the distinct necessary conditions, as recommended here. By contrast to this either/or, in a picture according to which conceptual contents are conferred on expressions by their being caught up in a structure of inferentially articulated commitments and entitlements, material inferential commitments are a necessary part of any package of practices that includes material doxastic commitments.

The circumstances and consequences of application of a nonlogical concept may stand in a substantive material-inferential relation. To ask what sort of 'harmony' they should exhibit is to ask what material inferences we ought to endorse, and so what conceptual contents we ought to employ. This is not the sort of a question to which we ought to expect or welcome a general or wholesale answer. Grooming our concepts and material-inferential commitments in the light of our assertional commitments, including those we find ourselves with noninferentially through observation, and the latter in the light of the former, is a messy, retail business.

Dummett thinks that a theory of meaning should take the form of an account of the nature of the 'harmony' that ought to obtain between the circumstances and the consequences of application of the concepts we ought to employ. Moving up a level now to apply these considerations about the relations of circumstances to consequences of application to the contents of the concepts employed in the metalanguage in which we couch a semantic theory, the point I want to make is that we should not expect a theory of that sort to take the form of a specification of necessary and sufficient conditions for the circumstances and consequences of application of a concept to be harmonious. For that presupposes that the circumstances and consequences of application of the concept of harmony do not themselves stand in a substantive material inferential relation. On the contrary, insofar as the idea of a theory of semantic or inferential harmony makes sense at all, it must take the form of an investigation of the ongoing elucidative process, of the 'Socratic method' of discovering and repairing discordant concepts, which alone gives the notion of harmony any content. It is given content only by the process of harmonizing commitments, from which it is abstracted. In Sellars' characterization of expressive rationality, modal claims are assigned the expressive role of inference licenses, which make explicit a commitment that is implicit in the use of conceptual contents antecedently in play. Rules of this sort assert an authority over future practice, and answer for their entitlement both to the prior practice being codified and to concomitant inferential and doxastic commitments. In this way

they may be likened to the principles formulated by judges at common law, intended both to codify prior practice, as represented by precedent, expressing explicitly as a rule what was implicit therein, and to have regulative authority for subsequent practice. The expressive task of making material inferential commitments explicit plays an essential role in the reflectively rational Socratic practice of harmonizing our commitments. For a commitment to become explicit is for it to be thrown into the game of giving and asking for reasons as something whose justification, in terms of other commitments and entitlements, is liable to question. Any theory of the sort of inferential harmony of commitments we are aiming at by engaging in this reflective, rational process must derive its credentials from its expressive adequacy to that practice, before it should be accorded any authority over it.

XIII: From Semantics To Pragmatics

In the first part of this essay I introduced three related ideas:

- the *inferential* understanding of conceptual content,
- the idea of *materially* good inferences, and
- the idea of *expressive* rationality.

These contrast, respectively, with:

--an understanding of content exclusively according to the model of the *representation* of states of affairs, [I think I've managed to say rather a lot about conceptual content here, without talking at all about what is represented by such contents.]

--an understanding of the goodness of inference exclusively on the model of *formal* validity, and

--an understanding of rationality exclusively on the model of *instrumental* or means-end reasoning.

In the second part of the essay, these ideas were considered in relation to the representation of inferential role suggested by Dummett, in terms of the circumstances of appropriate application of an expression or concept and the appropriate consequences of such application. It is in the context of these ideas that I have sought to present an *expressive* view of the role of logic, and its relation to the practices constitutive of rationality. That view holds out the hope of recovering for the study of *logic* a direct significance for projects that have been at the core of *philosophy* since its Socratic inception.

End

Panoramas of Mind and Meaning

Chapter Six:

How Analytic Philosophy has Failed Cognitive Science

1. We analytic philosophers have signally failed our colleagues in cognitive science. We have done that by not sharing central lessons about the nature of concepts, concept-use, and conceptual content that have been entrusted to our care and feeding for more than a century.

I take it that analytic philosophy began with the birth of the new logic that Gottlob Frege introduced in his seminal 1879 *Begriffsschrift*. The idea, taken up and championed to begin with by Bertrand Russell, was that the fundamental insights and tools Frege made available there, and developed and deployed through the 1890s, could be applied throughout philosophy to advance our understanding of understanding and of thought in general, by advancing our understanding of concepts—including the particular concepts with which the philosophical tradition had wrestled since its inception. For Frege brought about a revolution not just in *logic*, but in *semantics*. He made possible for the first time a *mathematical* characterization of meaning and conceptual content, and so of the structure of sapience itself. Henceforth it was to be the business of the new movement of analytic philosophy to explore and amplify those ideas, to exploit and apply

them wherever they could do the most good. Those ideas are the cultural birthright, heritage, and responsibility of analytic philosophers. But we have not done right by them. For we have failed to communicate some of the most basic of those ideas, failed to explain their significance, failed to make them available in forms usable by those working in allied disciplines who are also professionally concerned to understand the nature of thought, minds, and reason.

Contemporary cognitive science is a house with many mansions. The provinces I mean particularly to be addressing are cognitive psychology, developmental psychology, animal psychology (especially primatology), and artificial intelligence. (To be sure, this is not all of cognitive science. But the points I will be making in this paper are not of similarly immediate significance for such other subfields as neurophysiology, linguistics, perceptual psychology, learning theory, and the study of the mechanisms of memory.) Cognitive psychology aims at reverse-engineering the human mind: figuring out how we do what we do, what more basic abilities are recruited and deployed (and how) so as to result in the higher cognitive abilities we actually display. Developmental psychology investigates the sequence of stages by which those abilities emerge from more primitive versions as individual humans mature. Animal psychology, as I am construing it, is a sort of combination of cognitive psychology of non-human intelligences and a phylogenetic version of ontogenetic human developmental psychology. By contrast to all these empirical inquiries into actual cognition, artificial intelligence swings free of questions about how any actual organisms do what they do, and asks instead what constellation of abilities of the sort we know how to implement in artifacts might in principle yield sapience.

Each of these disciplines is in its own way concerned with the broadly empirical question of how the trick of cognition is or might be done. Philosophers are concerned with the normative question of what counts as doing it—with what understanding, particularly discursive, conceptual understanding consists in, rather than how creatures with a particular contingent constitution, history, and armamentarium of basic abilities come to exhibit it. I think Frege taught us three fundamental lessons about the structure of concepts, and hence about all possible abilities that deserve to count as concept-using abilities.⁷⁵ The conclusion we should draw from his discoveries is that concept-use is intrinsically stratified. It exhibits at least four basic layers, with each capacity to deploy concepts in a more sophisticated sense of ‘concept’ structurally presupposing the capacities to use concepts in all of the more primitive senses. The three lessons that generate the structural hierarchy oblige us to distinguish between:

- concepts that only *label* and concepts that *describe*,
- *ingredient* and *free-standing* conceptual contents, making explicit the distinction between the *content* of concepts and the *force* of applying them, and
- concepts expressible already by *simple* predicates and concepts expressible only by *complex* predicates.

AI researchers and cognitive, developmental, and animal psychologists need to take account of the different grades of conceptual content made visible by these distinctions, both in order to be clear about the topic they are investigating (if they are to tell us how the trick is done, they

⁷⁵ It ought to be uncontroversial that the last two of the three lessons are due to Frege. Whether he is responsible also for the first is more contentious. Further, I think both it and a version of the second can be found already in Kant. (As I argue in Part One of this book.) But my aims here are not principally hermeneutical or exegetical—those issues don’t affect the question of what we philosophers ought to be teaching cognitive scientists—so I will not be concerned to justify these attributions.

must be clear about exactly which trick it is) and because the empirical and in-principle possibilities are constrained by the way the abilities to deploy concepts in these various senses structurally presuppose the others that appear earlier in the sequence. This is a point they have long appreciated on the side of basic *syntactic* complexity. But the at least equally important—and I would argue more conceptually fundamental—hierarchy of *semantic* complexity has been largely ignored.

2. The Early Modern philosophical tradition was built around a *classificatory* theory of consciousness and (hence) of concepts, in part the result of what its scholastic predecessors had made of their central notion of Aristotelian forms. The paradigmatic cognitive act is understood as classifying: taking something particular *as* being of some general kind. Concepts are identified with those general kinds.

This conception was enshrined in the order of logical explanation (originating in Aristotle's *Prior Analytics*) that was common to *everyone* thinking about concepts and consciousness in the period leading up to Kant. At its base is a doctrine of *terms* or *concepts*, particular and general. The next layer, erected on that base, is a doctrine of *judgments*, describing the kinds of classificatory relations that are possible among such terms. For instance, besides classifying Socrates as human, humans can be classified as mortal. Finally, in terms of those metaclassifications grouping judgments into kinds according to the sorts of terms they relate, a doctrine of *consequences* or *syllogisms* is propounded, classifying valid inferences into

kinds, depending on which classes of classificatory judgments their premises and conclusions fall under.

It is the master-idea of classification that gives this traditional order of explanation its distinctive shape. That idea defines its base, the relation between its layers, and the theoretical aspiration that animates the whole line of thought: finding suitable ways of classifying terms and judgments (classifiers and classifications) so as to be able to classify inferences as good or bad solely in virtue of the kinds of classifications they involve. The fundamental metaconceptual role it plays in structuring philosophical thought about thought evidently made understanding the concept of classifying itself a particularly urgent philosophical task. Besides asking what differentiates various kinds of classifying, we can ask what they have in common. What is it one must *do* in order thereby to count as *classifying* something as being of some kind?

In the most general sense, one classifies something simply by responding to it differentially. Stimuli are grouped into kinds by the response-kinds they tend to elicit. In this sense, a chunk of iron classifies its environments into kinds by rusting in some of them and not others, increasing or decreasing its temperature, shattering or remaining intact. As is evident from this example, if classifying is just exercising a reliable differential responsive disposition, it is a ubiquitous feature of the inanimate world. For that very reason, classifying in this generic sense is not an attractive candidate for identification with conceptual, cognitive, or conscious activity. It doesn't draw the right line between thinking and all sorts of thoughtless activities. Pan-psychism is too high a price to pay for cognitive naturalism.

That need not mean that taking *differential responsiveness* as the genus of which *conceptual classification* is a species is a bad idea, however. A favorite idea of the classical British empiricists was to require that the classifying response be entering a *sentient* state. The intrinsic characters of these sentient states are supposed to sort them immediately into repeatable kinds. These are called on to function as the *particular* terms in the base level of the neo-Aristotelian logical hierarchy. *General* terms or concepts are then thought of as sentient state-kinds derived from the particular sentient state-kinds by a process of *abstraction*: grouping the base-level sentient state-repeatables into higher-level sentient state-repeatables by some sort of perceived *similarity*. This abstractive grouping by similarity is itself a kind of classification. The result is a path from one sort of consciousness, sentience, to a conception of another sort of consciousness, sapience, or conceptual consciousness.

A standing felt difficulty with this empiricist strategy is the problem of giving a suitably naturalistic account of the notion of sentient awareness on which it relies. Recent information-theoretic accounts of representation (under which heading I include not just Fred Dretske's theory, which actually goes by that name, but others such as Jerry Fodor's asymmetric counterfactual dependence and nomological locking models⁷⁶) develop the same basic differential responsiveness version of the classic classificatory idea in wholly naturalistic modal terms. They focus on the information conveyed about stimuli—the way they are grouped into repeatables—by their reliably eliciting a response of one rather than another repeatable response-

⁷⁶ Dretske, Fred: *Knowledge and the Flow of Information* (MIT Press—Bradford, 1981), Fodor, Jerry: *A Theory of Content* (MIT Press—Bradford, 1990). I discuss some aspects of such theories in Chapter Seven.

kind from some system. In this setting, unpalatable pan-psychism can be avoided not, as with traditional empiricism, by insisting that the responses be sentient states, but for instance by restricting attention to flexible systems, capable in principle of coming to encode many different groupings of stimuli, with a process of *learning* determining what classificatory dispositions each one actually acquires. (The classical American pragmatists' program for a naturalistic empiricism had at its core the idea that the structure common to evolutionary development and individual learning is a Test-Operate-Test-Exit negative feedback process of acquiring practical habits, including discriminative ones.⁷⁷)

Classification as the exercise of reliable differential responsive dispositions (however acquired) is not by itself yet a good candidate for *conceptual* classification, in the basic sense in which applying a concept to something is *describing* it. Why not? Suppose one were given a wand, and told that the light on the handle would go on if and only if what the wand was pointed at had the property of being *grivey*. One might then determine empirically that speakers are grivey, but microphones not, doorknobs are but windowshades are not, cats are and dogs are not, and so on. One is then in a position reliably, perhaps even infallibly, to apply the *label* 'grivey'. Is one also in a position to *describe* things *as* grivey? Ought what one is doing to qualify as applying the *concept* grivey to things? Intuitively, the trouble is that one does not know what one has found out when one has found out that something is grivey, does not know what one is

⁷⁷ I sketch this program in the opening section of "The Pragmatist Enlightenment (and its Problematic Semantics)" *European Journal of Philosophy*, Vol 12 No 1, April 2004, pp. 1-16.

taking it to be when one takes it to be grivey, does not know what one is describing it *as*. The label is, we want to say, uninformative.⁷⁸

What more is required? Wilfrid Sellars gives this succinct, and I believe correct, answer:

It is only because the expressions in terms of which we describe objects, even such basic expressions as words for the perceptible characteristics of molar objects, locate these objects in a space of implications, that they describe at all, rather than merely label.⁷⁹

The reason ‘grivey’ is merely a *label*, that it classifies without informing, is that nothing *follows* from so classifying an object. 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* 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

⁷⁸ This point is explored in greater detail in Chapter Seven.

⁷⁹ 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.

boxes to which they are applied. Description is classification with *consequences*, either immediately practical (“to be discarded/examined/kept”) or for further classifications.

Michael Dummett argues generally that to be understood as conceptually contentful, expressions must have not only *circumstances* of appropriate application, but also appropriate *consequences* of application.⁸⁰ That is, one must look not only *upstream*, to the circumstances (inferential and non-inferential) in which it is appropriate to apply the expression, but also *downstream* to the consequences (inferential and non-inferential) of doing so, in order to grasp the content it expresses. One-sided theories of meaning, which seize on one aspect to the exclusion of the other, are bound to be defective, for they omit aspects of the use that are essential to meaning. For instance, expressions can have the same circumstances of application, and different consequences of application. When they do, they will have different descriptive content.

1] I will write a book about Hegel,

and

2] I foresee that I will write a book about Hegel,

say different things about the world, describe it as being different ways. The first describes my future activity and accomplishment, the second my present aspiration. Yet the circumstances under which it is appropriate or warranted to assert them—the situations to which I ought

⁸⁰ I discuss this view of Dummett’s (from his *Frege: Philosophy of Language* second edition [Harvard University Press 1993], originally published in 1974), at greater length in Chapter Two of *Making It Explicit* [Harvard University Press, 1994], and Chapter One of *Articulating Reasons* [Harvard University Press, 2000].

reliably to respond by endorsing them—are the same (or at least, can be made so by light regimentation of a prediction-expressing use of ‘foresee’). Here, to say that they have different descriptive content can be put by saying that they have different truth conditions. (That they have the same assertibility conditions just shows how assertibility theories of meaning, as one-sided in Dummett’s sense, go wrong.) But that same fact shows up in the different positions they occupy in the “space of implications.” For from the former it follows that I will not be immediately struck by lightning, that I will write some book, and, indeed, that I will write a book about Hegel. None of these is in the same sense a consequence of the second claim.

We might train a parrot reliably to respond differentially to the visible presence of red things by squawking “That’s red.” It would not yet be *describing* things as red, would not be applying the concept red to them, because the noise it makes has no significance for it. It does not know that it follows from something’s being red that it is colored, that it cannot be wholly green, and so on. Ignorant as it is of those inferential consequences, the parrot does not grasp the concept (any more than we express a concept by ‘grivey’). The lesson is that even observational concepts, whose principal circumstances of appropriate application are non-inferential (a matter of reliable dispositions to respond differentially to non-linguistic stimuli) must have inferential consequences in order to make possible description, as opposed to the sort of classification effected by non-conceptual labels.

The rationalist idea that the inferential significance of a state or expression is essential to its *conceptual* contentfulness is one of the central insights of Frege’s 1879 *Begriffsschrift* (“concept

writing”)—the founding document of modern logic and semantics—and is appealed to by him in the opening paragraphs to define his topic:

...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...I call that part of the content that is the same in both the conceptual content [begriffliche Inhalt].⁸¹

Here, then, is the first lesson that analytic philosophy ought to have taught cognitive science: there is a fundamental meta-conceptual distinction between classification in the sense of *labeling* and classification in the sense of *describing*, and it consists in the *inferential* consequences of the classification: its capacity to serve as a premise in inferences (practical or theoretical) to further conclusions. (Indeed, there are descriptive concepts that are purely *theoretical*—such as gene and quark—in the sense that in addition to their inferential consequences of application, they have *only* inferential *circumstances* of application.) There is probably no point in fighting over the minimal circumstances of application of the concepts concept and conceptual. Those who wish to lower the bar sufficiently are welcome to consider purely classificatory labels as a kind of concept (perhaps so as not to be beastly to the beasts, or disqualify human infants, bits of our brains, or even some relatively complex computer programs wholly from engaging in conceptually articulated activities). But *if* they do so, they must *not*

⁸¹ Frege, *Begriffsschrift* (hereafter *BGS*), section 3. The passage continues: “In my formalized language [Begriffsschrift]...only that part of judgments which affects the possible inferences is taken into consideration. Whatever is needed for a correct inference is fully expressed; what is not needed is...not.”

combine those circumstances of application with the consequences of application appropriate to genuinely *descriptive* concepts—those that *do* come with inferential significances downstream from their application.

Notice that this distinction between labeling and describing is untouched by two sorts of elaborations of the notion of labeling that have often been taken to be of great significance in thinking about concepts from the classical classificatory point of view. One does not cross the boundary from labeling to describing just because the reliable capacity to respond differentially is *learned*, and in that sense flexible, rather than *innate*, and in that sense rigid. And one is likewise developing the classical model in an orthogonal direction insofar as one focuses on the metacapacity to learn to distinguish arbitrary Boolean combinations of microfeatures one can already reliably discriminate. From the point of view of the distinction between labeling and describing, that is not yet the capacity to form *concepts*, but only the mastery of *compound* labels. That sort of structural articulation upstream has no *semantic* import at the level of description until and unless it is accorded a corresponding inferential significance downstream.

3. Once our attention has been directed at the significance of applying a classifying concept—downstream, at the consequences of applying it, rather than just upstream, at the repeatable it discriminates, the grouping it institutes—so that *mere* classification is properly distinguished from *descriptive* classification, the necessity of distinguishing different *kinds* of consequence becomes apparent. One distinction in the vicinity, which has already been mentioned in passing, is that between *practical* and *theoretical* (or, better, *cognitive*)

consequences of application of a concept. The significance of classifying an object by responding to it one way rather than another may be to make it appropriate to *do* something else with or to it—to keep it, examine it, or throw it away, to flee or pursue or consume it, for example. This is still a matter of inference; in this case, it is *practical* inferences that are at issue. But an initial classification may also contribute to further classifications: that what is in my hand falls under both the classifications raspberry and red makes it appropriate to classify it also as ripe—which in turn has practical consequences of application (such as, under the right circumstances “falling to without further ado and eating it up,” as Hegel says in another connection) that neither of the other classifications has individually. Important as the distinction between practical and cognitive inferential consequences is, in the present context there is reason to emphasize a different one.

Discursive intentional phenomena (and their associated concepts), such as assertion, inference, judgment, experience, representation, perception, action, endorsement, and imagination typically involve what Sellars calls “the notorious ‘ing’/‘ed’ ambiguity.” For under these headings we may be talking about the *act* of asserting, inferring, judging, experiencing, representing, perceiving, doing, endorsing, and imagining, or we may be talking about the *content* that is asserted, inferred, judged, experienced, represented, perceived, done, endorsed, or imagined. ‘Description’ is one of these ambiguous terms (as is ‘classification’). We ought to be aware of the distinction between the act of describing (or classifying), applying a concept, on the one hand, and the content of the description (classification, concept)—*how* things are described (classified, conceived)—on the other. And the distinction is not merely of theoretical importance for those of us thinking systematically about concept use. A distinctive level of conceptual

sophistication is achieved by concept users that themselves distinguish between the contents of their concepts and their activity of applying them. So one thing we might want to know about a system being studied, a non-human animal, a prelinguistic human, an artifact we are building, is whether *it* distinguishes between the *concept* it applies and what it *does* by applying it.

We can see a basic version of the distinction between semantic content and pragmatic force as in play wherever *different* kinds of practical significance can be invested in the *same* descriptive content (different sorts of speech act or mental act performed using that content). Thus if a creature can not only say or think that the door is shut, but also ask or wonder whether the door is shut, or order or request that it be shut, we can see it as distinguishing in practice between the content being expressed and the pragmatic force being attached to it. In effect, it can use descriptive contents to do more than merely describe. But this sort of practical distinguishing of pragmatic from semantic components matters for the semantic hierarchy I am describing only when it is incorporated or reflected in the *concepts* (that is, the *contents*) a creature can deploy. The capacity to attach different sorts of pragmatic force to the same semantic content is not sufficient for *this* advance in structural semantic complexity. (Whether it is a necessary condition is a question I will not address—though I am inclined to think that in principle the answer is ‘No’.)

For the inferential consequences of applying a classificatory concept, when doing that is describing and not merely labeling, can be either *semantic* consequences, which turn on the *content* of the concept being applied, or *pragmatic* consequences, which turn on the *act* one is

performing in applying it. Suppose John issues an observation report: “The traffic light is red.” You may infer that it is operating and illuminated, and that traffic ought to stop in the direction it governs. You may also infer that John has a visually unobstructed line of sight to the light, notices what color it is, and believes that it is red. Unlike the former inferences, these are not inferences from what John *said*, from the *content* of his utterance, from the concepts he has applied. They are inferences from his *saying* it, from the pragmatic force or significance of his *uttering* it, from the fact of his *applying* those concepts. For what he has *said*, that the traffic light is red, could be true even if John had not been in a position to notice it or form any beliefs about it. Nothing about John follows just from the color of the traffic light.⁸²

It can be controversial whether a particular consequence follows from how something is described or from describing it that way, that is, whether that consequence is part of the descriptive content of an expression, the concept applied, or stems rather from the force of using the expression, from applying the concept. A famous example is expressivist theories of evaluative terms such as ‘good’. In their most extreme form, they claim that these terms have no descriptive content. *All* their consequences stem from what one is doing in using them: commending, endorsing, or approving. In his lapidary article “Ascriptivism,”⁸³ Peter Geach asks what the rules governing this move are. He offers the archaic term ‘macarize’, meaning to

⁸² One might think that a similar distinction could be made concerning a parrot that merely reliably responsively discriminated red things by squawking “That’s red.” For when he does that, one might infer that there was something red there (since he is reliable), and one might also infer that the light was good and his line of sight unobstructed. So both sorts of inference seem possible in this case. But it would be a mistake to describe the situation in these terms. The squawk is a label, not a description. We infer from the parrot’s producing it that there is something red, because the two sorts of events are reliably correlated, just as we would from the activation of a photocell tuned to detect the right electromagnetic frequencies. By contrast, John offers *testimony*. What he says is usable as a premise in our own inferences, not just the fact that his saying it is reliably correlated with the situation he (but not the parrot) reports (though they both respond to it).

⁸³ *The Philosophical Review*, Vol. 69, No. 2, 221-225. Apr., 1960.

characterize someone as happy. Should we say that in apparently describing someone as happy we are not really describing anyone, but rather performing the distinctive speech act of macarizing? But why not then discern a distinctive speech act for *any* apparently descriptive term?

What is wanted is a criterion for distinguishing semantic from pragmatic consequences, those that stem from the content of the concept being applied from those that stem from what we are doing in applying that concept (using an expression to perform a speech act). Geach finds one in Frege, who in turn was developing a point made already by Kant.⁸⁴ The logical tradition Kant inherited was built around the classificatory theory of consciousness we began by considering. Judgment was understood as classification or predication: paradigmatically, *of* something particular *as* something general. But we have put ourselves in a position to ask: is this intended as a model of how judgeable contents are constructed, or of what one is doing in judging? Kant saw, as Frege would see after him, that the phenomenon of *compound* judgments shows that it *cannot* play both roles. For consider the hypothetical or conditional judgment:

3] If Frege is correct, then conceptual content depends on inferential consequences.

In asserting this sentence (endorsing its content), have I predicated correctness of Frege (classified him as correct)? Have I described him as correct? Have I applied the concept of correctness? If so, then predicating or classifying (or describing) is not judging. For in asserting the conditional I have *not* judged or asserted that Frege is correct. I have at most built up a judgeable content, the antecedent of the conditional, by predication. For embedding a

⁸⁴ I discuss this point further in Chapter One.

declarative descriptive sentence as an unasserted component in a compound asserted sentence strips off the pragmatic force its free-standing, unembedded occurrence would otherwise have had. It now contributes *only* its *content* to the *content* of the compound sentence, to which alone the pragmatic force of a speech act is attached.⁸⁵

This means that embedding simpler sentences as components of compound sentences—paradigmatically, embedding them as antecedents of conditionals—is the way to discriminate consequences that derive from the *content* of a sentence from consequences that derive from the *act* of asserting or endorsing it. We can tell that ‘happy’ *does* express descriptive content, and is *not* simply an indicator that some utterance has the pragmatic force or significance of macarizing, because we *can* say things like:

4] If she is happy, then John should be glad.

For in asserting that, one does *not* macarize anyone. So the consequence, that John should be glad, must be due to the descriptive content of the antecedent, not to its force.

Similarly, Geach argues that the fact that we can say things like:

5] If being trustworthy is good, then you have reason to be trustworthy,

shows that ‘good’ *does* have descriptive content.⁸⁶ Notice that this same test appropriately discriminates the different descriptive contents of the claims:

⁸⁵ This point is further expounded and exploited in Chapter One.

⁸⁶ Of course, contemporary expressivists such as Gibbard and Blackburn (who are distinguished from emotivist predecessors such as C.L. Stevenson precisely by their appreciation of the force of the Frege-Geach argument) argue that it need not follow that the right way to understand that descriptive content is not by tracing it back to the

6] Labeling is not describing,

and

7] I believe that labeling is not describing.

For the two do not behave the same way as antecedents of conditionals. The stuttering inference

8] If labeling is not describing, then labeling is not describing,

is as solid an inference as one could ask for. The corresponding conditional

9] If I believe that labeling is not describing, then labeling is not describing,

requires a good deal more faith to endorse. And in the same way, the embedding test

distinguishes [1] and [2] above. In each case it tells us, properly, that different descriptive contents are involved.

What all this means is that any user of descriptive concepts who can also form compound sentences, paradigmatically conditionals, is in a position to distinguish what pertains to the semantic *content* of those descriptive concepts from what pertains to the *act* or pragmatic *force* of describing by applying those concepts. This capacity is a new, higher, more sophisticated level of concept use. It can be achieved *only* by looking at compound sentences in which other descriptive sentences can occur as unasserted components. For instance, it is only in such a context that one can distinguish *denial* (a kind of speech act or attitude) from *negation* (a

attitudes of endorsement or approval that *are* expressed by the use of the expression in free-standing, unembedded assertions.

kind of content). One who asserts [6] has *both* denied that labeling is describing, *and* negated a description. But one who asserts conditionals such as [8] and [9] has negated descriptions, but has *not* denied anything.

The modern philosophical tradition up to Frege took it for granted that there was a special attitude one could adopt towards a descriptive conceptual content, a kind of minimal force one could invest it with, that must be possible independently of and antecedently to being able to endorse that content in a judgment. This is the attitude of merely *entertaining* the description. The picture (for instance, in Descartes) was that *first* one entertained descriptive thoughts (judgeables), and *then*, by an in-principle subsequent act of will, accepted or rejected it. Frege rejects this picture. The principal—and in principle fundamental—pragmatic attitude (and hence speech act) is judging or endorsing.⁸⁷ The capacity merely to entertain a proposition (judgeable content, description) is a late-coming capacity—one that is parasitic on the capacity to endorse such contents. In fact, for Frege, the capacity to entertain (without endorsement) the proposition that *p* is just the capacity to endorse *conditionals* in which that proposition occurs as antecedent or consequent. For that is to explore its descriptive content, its inferential circumstances and consequences of application, what it follows from and what follows from it, what would make it true and what would be true if it were true, without endorsing it. This is a new kind of distanced attitude toward one's concepts and their contents—one that becomes possible only in virtue of the capacity to form compound sentences of the kind of which conditionals are the paradigm. It is a new level of cognitive achievement—not in the sense of a new kind of empirical knowledge

⁸⁷ In Chapter One I discuss the line of thought that led Kant to give pride of place to judgment and judging.

(though conditionals can indeed codify new empirical discoveries), but of a new kind of semantic self-consciousness.

Conditionals make possible a new sort of hypothetical thought. (Supposing that postulating a distinct attitude of supposing would enable one to do *this* work, the work of conditionals, would be making the same mistake as thinking that denial can do the work of negation.) Descriptive concepts bring empirical properties into view. Embedding those concepts in conditionals brings the contents of those concepts into view. Creatures that can do that are functioning at a higher cognitive and conceptual level than those who can only apply descriptive concepts, just as those who can do that are functioning at a higher cognitive and conceptual level than those who can only classify things by reliable responsive discrimination (that is, labeling). That fact sets a question for the different branches of cognitive science I mentioned in my introduction. Can chimps, or African grey parrots, or other non-human animals not just use concepts to describe things, but also semantically discriminate the contents of those concepts from the force of applying them, by using them not just in describing, but in conditionals, in which their contents are merely entertained and explored? At what age, and along with what other capacities, do human children learn to do so? What is required for a computer to demonstrate this level of cognitive functioning?

Conditionals are special, because they make *inferences* explicit—that is, put them into endorsable, judgeable, assertible, which is to say propositional form. And it is their role in inferences, we saw, that distinguishes descriptive concepts from mere classifying labels. But

conditionals are an instance of a more general phenomenon. For we can think of them as operators, which apply to sentences to yield further sentences. As such, they bring into view a new notion of conceptual content: a new principle of assimilation, hence classification, of such contents. For we begin with the idea of sameness of content that derives from sameness of pragmatic force, attitude, or speech act. But the Frege-Geach argument shows that we can also individuate conceptual contents more finely, not just in terms of their role in free-standing utterances, but also accordingly as substituting one for another as arguments of operators (paradigmatically the conditional) does or does not yield compound sentences with the same free-standing pragmatic significance or force. Dummett calls these notions “free-standing” and “ingredient” content (or sense), respectively. Thus we might think that

10] It is nice here,

and

11] It is nice where I am,

express the same attitude, perform the same speech act, have the same pragmatic force or significance. They not only have the same circumstances of application, but the same consequences of application (and hence role as antecedents of conditionals). But we can see that they have different *ingredient* contents by seeing that they behave differently as arguments when we apply another operator to them. To use an example of Dummett’s,

12] It is *always* nice here,

and

13] It is *always* nice where I am,

have very different circumstances and consequences of application, different pragmatic significances, and *do* behave differently as the antecedents of conditionals. But this difference in content, this sense of “different content” in which they patently do have different contents, is one that shows up *only* in the context of compounding operators, which apply to sentences and yield further sentences. The capacity to deploy such operators to form new conceptual (descriptive) contents from old ones accordingly ushers in a new level of cognitive and conceptual functioning.

Creatures that can not merely label, but describe are *rational*, in the minimal sense that they are able to treat one classification as providing a *reason* for or against another. If they can use conditionals, they can distinguish inferences that depend on the *content* of the concept they are applying from those that depend on what they are *doing* in classifying something as falling under that concept. But the capacity to use conditionals gives them more than just that ability. For conditionals let them *say* what is a reason for what, say *that* an inference is a good one. And for anyone who can do that, the capacity not just to *deny* that a classification is appropriate, but to use a *negation* operator to form new classificatory contents brings with it the capacity to say that two classifications (classifiers, concepts) are incompatible: that one provides a reason to withhold the other. Creatures that can use this sort of sentential compounding operator are not just *rational*, but *logical* creatures. They are capable of a distinctive kind of *conceptual self-consciousness*. For they can describe the rational relations that make their classifications into descriptions in the first place, hence be conscious or aware of them in the sense in which descriptive concepts allow them to be aware of empirical features of their world.

4. There is still a higher level of structural complexity of concepts and concept use. I have claimed that Frege should be credited with appreciating both of the points I have made so far: that descriptive conceptual classification beyond mere discriminative labeling depends on the inferential significance of the concepts, and that semantically distinguishing the inferential significance of the contents of concepts from that of the force of applying them depends on forming sentential compounds (paradigmatically conditionals) in which other sentences appear as components. In each of these insights Frege had predecessors. Leibniz (in his *New Essay on the Human Understanding*) had already argued the first point, against Locke. (The move from thinking of concepts exclusively as reliably differentially elicited labels to thinking of them as having to stand in the sort of inferential relations to one another necessary for them to have genuine descriptive content is characteristic of the advance from empiricism to rationalism.) And Kant, we have seen, appreciated how attention to compound sentences (including “hypotheticals”) requires substantially amending the traditional classificatory theory of conceptual consciousness. The final distinction I will discuss, that between *simple* and *complex* predicates, and the corresponding kinds of concepts they express, is Frege’s alone. No-one before him (and embarrassingly few even of his admirers after him) grasped this idea.

Frege’s most famous achievement is transforming traditional logic by giving us a systematic way to express and control the inferential roles of *quantificationally complex* sentences. Frege could, as the whole logical tradition from Aristotle down to his time (fixated as

it was on syllogisms) could not, handle iterated quantifiers. So he could, for instance, explain why

14] If someone is loved by everyone, then everyone loves someone,

is true (a conditional that codifies a correct inference), but

15] If everyone loves someone, then someone is loved by everyone,

is not. What is less appreciated is that in order to specify the inferences involving arbitrarily nested quantifiers ('some' and 'every'), he needed to introduce a new kind of predicate, and hence discern a structurally new kind of *concept*.

Our first grip on the notion of a predicate is as a *component* of sentences. In artificial languages we combine, for instance, a two-place predicate 'P' with two individual constants 'a' and 'b' to form the sentence 'Pab'. Logically minded philosophers of language use this model to think about the corresponding sentences of natural languages, understanding

16] Kant admired Rousseau,

as formed by applying the two-place predicate 'admired' to the singular terms 'Kant' and 'Rousseau'. The kind of inferences that are made explicit by *quantified conditionals*—inferences that essentially depend on the contents of the predicates involved—though, require us also to distinguish a one-place predicate, related to but distinct from this two-place one, that is exhibited by

17] Rousseau admired Rousseau,

and

18] Kant admired Kant,

but *not* by [16].

19] Someone admired himself,

that is, something of the form $\exists x[Pxx]$, follows from [17] and [18], but not from [16]. The property of **being a self-admirer** differs from that of **being an admirer** and from that of **being admired** (even though it entails both).

But there is no *part* of the sentences [17] and [18] that they share with each other that they do *not* share also with [16]. Looking just at the sub-sentential expressions out of which the sentences are built does not reveal the respect of similarity that distinguishes self-admiration from admiration in general—a respect of similarity that is crucial to understanding why the conditional

20] If someone admires himself then someone admires someone,

$(\exists x[Pxx] \rightarrow \exists x \exists y [Pxy])$ expresses a good inference, while

21] If someone admires someone then someone admires himself,

$(\exists x\exists y[Pxy] \rightarrow \exists x[Pxx])$ does not. For what [17] and [18] share that distinguishes them from [16] is not a *component*, but a *pattern*. More specifically, it is a pattern of cross-identification of the singular terms that two-place predicate applies to.

The repeatable expression-kind ‘admires’ is a *simple* predicate. It occurs as a component in sentences built up by concatenating it appropriately with a pair of singular terms. ‘x admires x’ is a *complex* predicate.⁸⁸ A number of different complex predicates are associated with any multi-place simple predicate. So the three-place simple predicate used to form the sentence

22] John enjoys music recorded by Mark and books recommended by Bob,

generates not only a three-place complex predicate of the form $Rxyz$, but also two-place complex predicates of the form $Rxxy$, $Rxyy$, and $Rxyx$, as well as the one-place complex predicate $Rxxx$. The complex predicates can be thought of as patterns that can be exhibited by sentences formed using the simple predicate, or as equivalence classes of such sentences. Thus the complex self-admiration predicate can be thought of either as the *pattern*, rather than the *part*, that is common to all the sentences {“Rousseau admired Rousseau,” “Kant admired Kant,” “Caesar admired Caesar,” “Brutus admired Brutus,” “Napoleon admired Napoleon,”...}, or just as that set itself. Any member of such an equivalence class of sentences sharing a complex predicate can be turned into any other by a sequence of *substitutions* of all occurrences of one singular term by occurrences of another.

⁸⁸ This point, and the terminology of ‘simple’ and ‘complex’ predicates, is due to Dummett, in the second chapter of his monumental *Frege’s Philosophy of Language* [op.cit.].

Substitution is a kind of *decomposition* of sentences (including compound ones formed using sentential operators such as conditionals). After sentences have been built up using simple components (singular terms, simple predicates, sentential operators), they can be assembled into equivalence classes (patterns can be discerned among them) by regarding some of the elements as systematically replaceable by others. This is the same procedure of noting invariance under substitution that we saw applies to the notion of free-standing content to give rise to that of ingredient content, when the operators apply only to whole sentences. Frege called what is invariant under substitution of some sentential components for others a '*function*'. A function can be applied to some arguments to yield a value, but it is not a *part* of the value it yields. (One can apply the function *capital of* to Sweden to yield the value Stockholm, but neither Sweden nor *capital of* is part of Stockholm.) He tied himself in some metaphysical knots trying to find a clear way of contrasting functions with *things* (objects). But two points emerge clearly. First, discerning the substitutional relations among different sentences sharing the same simple predicate is crucial for characterizing a wide range of inferential patterns. Second, those inferential patterns articulate the contents of a whole new class of concepts.

Sentential compounding already provided the means to build new simple concepts out of old ones. The Boolean connectives—conjunction, disjunction, negation, and the conditional definable in terms of them ($A \rightarrow B$ if and only if $\sim(A \& \sim B)$)—permit the combination of simple predicates in all the ways representable by Venn diagrams, corresponding to the intersection, union, complementation, and inclusion of sets (concept extensions, represented by regions), and

so the expression of new concepts formed from old ones by these operations. But there is a crucial class of new analytically *complex* concepts formable from the old ones that are *not* generable by such compounding procedures. One cannot, for instance, form the concept of a C such that for every A there is a B that stands to that C in the relation R. This is the complex one-place predicate logicians would represent as having the form $\{x: Cx \ \& \ \forall y \in A \exists z \in B [Ryz]\}$. As Frege says, such a concept cannot, as the Boolean ones can, be formed simply by putting together pieces of the boundaries of the concepts A, B, and C. The correlations of elements of these sets that concepts like these, those expressed by complex predicates, depend on, and so the inferences they are involved in, cannot be represented in Venn diagrams.

Frege showed further that it is just concepts like these that even the simplest mathematics works with. The concept of a natural number is the concept of a set every element of which has a successor. That is, for every number, there is another related to it as a successor ($\forall x \exists y [\text{Successor}(x, y)]$). The decisive advance that Frege's new quantificational logic made over traditional logic is a *semantic, expressive* advance. His logical notation can, as the traditional logic could not, form *complex* predicates, and so both express a vitally important kind of concept, and logically codify the inferences that articulate its descriptive content.

Complex concepts can be thought of as formed by a four-stage process.

- First, put together simple predicates and singular terms, to form a set of sentences, say $\{Rab, Sbc, Tacd\}$.

- Then apply sentential operators to form compound sentences, say $\{Rab \rightarrow Sbc, Sbc \& Tacd\}$.
- Then substitute variables for some of the singular terms (individual constants), to form complex predicates, say $\{Rax \rightarrow Sxy, Sxy \& Tayz\}$.
- Finally, apply quantifiers to bind some of these variables, to form new complex predicates, for instance the one-place predicates (in y and z) $\{\exists x[Rax \rightarrow Sxy], \forall x \exists y[Sxy \& Tayz]\}$.

If one likes, this process can now be repeated, with the complex predicates just formed playing the role that simple predicates originally played at the first stage, yielding the new sentences $\{\exists x[Rax \rightarrow Sxd], \forall x \exists y[Sxy \& Taya]\}$. They can then be conjoined, and the individual constant a substituted for to yield the further one-place complex predicate (in z) $\exists x[Rzx \rightarrow Sxd] \& \forall x \exists y[Sxy \& Tzyz]$. We can use these procedures to build to the sky, repeating these stages of concept construction as often as we like. Frege's rules tell us how to compute the inferential roles of the concepts formed at each stage, on the basis of the inferential roles of the raw materials, and the operations applied at that stage. This is the heaven of complex concept formation he opened up for us.

5. The result of all these considerations, which have been in play since the dawn of analytic philosophy, well over a century ago, is a four-stage *semantic* hierarchy of ever more demanding senses of “concept” and “concept use.” At the bottom are concepts as reliably differentially applied, possibly learned, *labels* or classifications. Crudely behaviorist psychological theories (such as B. F. Skinner's) attempted to do all their explanatory work with responsive

discriminations of this sort. At the next level, concepts as *descriptions* emerge when merely classifying concepts come to stand in *inferential, evidential, justificatory* relations to one another—when the propriety of one sort of classification has the practical significance of making others appropriate or inappropriate, in the sense of serving as *reasons* for and against them. Concepts of this sort may still all have observational uses, even though they are distinguished from labels by also having inferential ones.⁸⁹ Already at this level, the possibility exists of empirical descriptive concepts that can *only* be properly applied as the result of inferences from the applicability of others. These are *theoretical* concepts: a particularly sophisticated species of the genus of descriptive concepts.

At this second level, conceptual content first takes a distinctive *propositional* form; applications of this sort of concept are accordingly appropriately expressed using *declarative sentences*. For the propositional contents such sentences express just are whatever can play the role of premise and conclusion in *inferences*. And it is precisely being able to play those roles that distinguishes applications of descriptive concepts from applications of merely classificatory ones. Building on the capacity to use inferentially articulated descriptive concepts to make propositionally contentful judgments or claims, the capacity to form sentential *compounds*—paradigmatically *conditionals*, which make endorsements of material inferences relating descriptive concept applications propositionally explicit, and *negations*, which make endorsements of material incompatibilities relating descriptive concept applications

⁸⁹ A key part of the higher *inferential* grade of conceptuality (which includes the former, but transforms it) is that it is *multipremise material* inferences that one learns to draw as conclusions (=responses) now to Boolean combinations of the relatively enduring states that result from one's own responses.

propositionally explicit—brings with it the capacity to deploy a further, more sophisticated, kind of conceptual content: *ingredient* (as opposed to free-standing) content. Conceptual content of this sort is to be understood in terms of the contribution it makes to the content of *compound* judgments in which it occurs, and only thereby, indirectly, to the force or pragmatic significance of endorsing that content.

Ingredient conceptual content, then, is what can be *negated*, or *conditionalized*. The distinctive sort of definiteness and determinateness characteristic of this sort of conceptual content becomes vivid when it is contrasted with contents that cannot appear in such sentential compounds, such as that expressed by pictures. My young son once complained about a park sign consisting of the silhouette of what looked like a Scottish terrier, surrounded by a red circle, with a slash through it. Familiar with the force of prohibition associated with signs of this general form, he wanted to know: “Does this mean ‘No Scotties allowed’? Or ‘No dogs allowed’? Or ‘No animals allowed’? Or ‘No pets allowed’”? Indeed. With pictures one has no way of indicating the degree of generality intended. A creature that can understand a claim like “If the red light is on, then there is a biscuit in the drawer,” without disagreeing when the light is not on and no biscuit is present, or immediately looking for the biscuit regardless of how it is with the light, has learned to distinguish between the content of descriptive concepts and the force of applying them, and as a result can entertain and explore those concepts and their connections with each other without necessarily applying them in the sense of endorsing their applicability to anything present. The capacity in this way to free oneself from the bonds of the here-and-now is a distinctive kind of conceptual achievement

The first step was from merely *discriminating* classification to *rational* classification ('rational' because inferentially articulated, according to which classifications provide reasons for others). The second step is to *synthetic logical* concept formation, in which concepts are formed by logical compounding operators, paradigmatically conditionals and negation. The final step is to *analytical* concept formation, in which the sentential compounds formed at the third stage are *decomposed* by noting invariants under substitution. This is actually the same method that gave us the notion of ingredient content at the third stage of concept formation. For that metaconcept arises when we realize that two sentences that have the same pragmatic potential as free-standing, force-bearing rational classifications can nonetheless make different contributions to the content (and hence the force) of compound sentences in which they occur as unendorsed components—that is, when we notice that substituting one for the other may change the free-standing significance of asserting the compound sentence containing them. To form *complex* concepts, we must apply the same methodology to sub-sentential expressions, paradigmatically singular terms, that have multiple occurrences in those same logically *compound* sentences. Systematically assimilating sentences into various equivalence classes accordingly as they can be regarded as substitutional variants of one another is a distinctive kind of *analysis* of those compound sentences, as involving the application of concepts that were not *components* out of which they were originally constructed. Concepts formed by this sort of analysis are substantially and in principle more expressively powerful than those available at earlier stages in the hierarchy of conceptual complexity. (They are, for instance, indispensable for even the simplest mathematics.)

This hierarchy is not a *psychological* one, but a *logical* and *semantic* one. Concepts at the higher levels of complexity presuppose those at lower levels, not because creatures of a certain kind cannot in practice, as a matter of fact, deploy the more complex kinds unless they can deploy the simpler ones, but because in principle it is structurally impossible to do so. Nothing could count as grasping or deploying the kinds of concepts that populate the upper reaches of the hierarchy without also grasping or deploying those drawn from its lower levels. The dependencies involved are not empirical, but (meta)conceptual and normative. The Fregean considerations that enforce the distinctions between and sequential arrangement of concept-kinds do not arise from studying how concept-users actually work, but from investigation of what concept use fundamentally is. They concern not how the trick (of concept use) is done, but what could in principle count as doing it—a normative, rather than an empirical issue. That is why it is philosophers who first came across this semantic hierarchical metaconceptual structure of concept-kinds.

But cognitive scientists need to know about it. For it is part of the job of the disciplines that cognitive science comprises to examine—each from its own distinctive point of view—all four grades of conceptual activity: the use of more complex and sophisticated kinds of concepts, no less than that of the simpler and less articulated sorts. The move from merely classificatory to genuinely descriptive concepts, for instance, marks a giant step forward in the phylogenetic development of sapience. I do not think we yet know what non-human creatures are capable of taking that step. Human children clearly do cross that boundary, but when, by what means, at what age or stage of development? Can non-human primates learn to use conditionals? Has anyone ever tried to teach them? The only reason to focus on that capacity, out of all the many

linguistic constructions one might investigate empirically in this regard, is an appreciation of the kind of semantic self-consciousness about the rational relations among classifications (which marks the move from classification to rational description) that they make possible. Computer scientists have, to be sure, expended some significant effort in thinking about varieties of possible implementation of sentential compounding—for instance in exploring what connectionist or parallel distributed processing systems can do. But they have not in the same way appreciated the significance of the question of whether, to what extent, and how such “vehicleless” representational architectures can capture the full range of concepts expressed by complex predicates. (Those systems’ lack of syntactically compositional explicit symbolic representations prohibits the standard way of expressing these concepts, for that way proceeds precisely by substitutional *decomposition* of such explicit symbolic representations.) These are merely examples of potentially important questions raised by the hierarchy of conceptual complexity that cognitive scientists have by and large not been moved so much as to ask.

Why not? I think it is pretty clear that the answer is *ignorance*. Specifically, it is ignorance of the considerations, put forward already by Frege, that draw the bright semantic metaconceptual lines between different grades of concepts, and arrange them in a strict presuppositional semantic hierarchy. Any adequately trained cognitive scientist—even those working in disciplines far removed from computational linguistics—can be presumed to have at least passing familiarity with the similarly four-membered Chomsky hierarchy that lines up kinds of grammar, automaton, and *syntactic* complexity of languages in an array from most basic (finite state automata computing regular languages specifiable by the simplest sort of grammatical rules) to most sophisticated (two-stack pushdown automata computing recursively

enumerable languages specifiable by unrestricted grammatical rules). But the at least equally significant *semantic* distinctions I have been retailing have not similarly become a part of the common wisdom and theoretical toolbox of cognitive science—even though they have been available for a half-century longer.

The cost of that ignorance, in questions not asked, theoretical constraints not appreciated, promising avenues of empirical research not pursued, is great. Failure to appreciate the distinctions and relations among fundamentally different kinds of concepts has led, I think, to a standing tendency systematically to overestimate the extent to which one has constructed (in AI) or discerned in development (whether by human children or non-human primates) or reverse-engineered (in psychology) what *we* users of the fanciest sorts of concepts do. That underlying ignorance is culpable. But it is not the cognitive scientists themselves who are culpable for their ignorance. The ideas in question are those that originally launched the whole enterprise of analytic philosophy. I think it is fair to say that as we philosophers have explored these ideas, we have gotten clearer about them in many respects. For one reason or another, though, we have not shared the insights we have achieved. We are guilty of having kept this treasure trove to ourselves. It is high time to be more generous in sharing these ideas.

End

Panoramas of Mind and Meaning

Chapter Seven:

Reason, Genealogy, and the Hermeneutics of Magnanimity

I. A Metanarrative: from Disenchantment to Disillusionment

Hegel said: “To him who looks on the world rationally, the world looks rationally back.”⁹⁰ More than half a century later, Nietzsche said “When you stare long into the abyss, the abyss stares back into you.”⁹¹ These paired Spiegeleiern passages express, in gnomic aphorisms, sentiments that mark the endpoints of a critical arc of nineteenth century philosophical thought. Hegel’s sunny homily epitomizes the optimism of his version of the Enlightenment rationalism that flourished in the previous century. Nietzsche’s darker remark foreshadows the pessimism of a distinctive kind of nihilism, rooted in reductive naturalism, which the events of the following century would make both familiar and fitting. Each of these successive nineteenth century currents of thought, one looking back to what had already been understood and one pointing

⁹⁰ Introduction to *Lectures on the Philosophy of History*, published separately as *Reason in History*, tr. R.S. Hartman 1953, p. 13; Suhrkamp Hegel edition, volume 12, p.23.

⁹¹ From *Beyond Good and Evil*, Aphorism 146.

ahead to what had yet to be dealt with, comes with a rationalizing narrative of *progress*: the first, of disenchantment *by* reason, the second of disillusionment *with* reason.

It was always essential to the self-understanding of Enlightenment that it see itself as the advent of something both genuinely new and essentially progressive. It defined itself by the contrast between the light of reason it sought, developed, and celebrated, and the darkness from which Enlightenment arose, and by which it was still surrounded and would always be threatened: the shadows of superstition, prejudice, and dogmatism cast by arbitrary despotic power sedimented in the merely traditional institutions with which those habits of thought connived and in which they thrived.

The fundamental conceptual innovation of the time was not the focus on reason by itself. Philosophy, whose avatar is Socrates, had perennially championed reason. Nor is it the mere association of reason with *freedom*. “Know the truth, and the truth shall set you free,” the Christian tradition, in the person of John, already taught.. What is wholly new in Enlightenment philosophy, its characteristic insight, is its identification of that transformative emancipatory power with reason in its *critical* function. The *only* authority it admits as legitimate and legitimating is the authority of the better reason—that peculiar normative force, compelling only to the rational, that had so fascinated and puzzled the Greeks. And the Enlightenment acknowledges no higher judge competent to assess the merits of competing reasons than the “natural light” with which the capacities of each individual reasoning subject equip it. That is why Kant says “Sapere aude!” (“Dare to understand”)—that is the motto of enlightenment,” in

his essay identifying enlightenment as “man’s release from his self-imposed tutelage.”⁹² The advent of an age in which individuals accept no authority transcending their own capacity critically to assess reasons is for Kant, speaking here for the whole Enlightenment, nothing less than humanity’s coming to maturity. This emancipation (literally: in Roman law, the process by which children are set free from the *patria potestas*) is to be effected by wholesale replacement of the traditional model of authority, which understands it exclusively in terms of the *obedience* owed by a subordinate to a superior (metaphysically requiring an underlying metaphysical normatively determinative *scala naturae*: the Great Chain of Being), by a model that understands authority exclusively in terms of the force of impersonal reasons, assessable by all. Reason, for Kant, can accordingly be identified as *freedom* in the form of *autonomy*. The authority of the superior-in-power is abolished. Authority resides only in one’s own acknowledgement of reasons, which are reasons for all alike.

All the great philosophers in the period from Descartes to Kant were theorists of Enlightenment. Hegel, though, is the first to take the advent of *modernity*—for him, the single most important thing that has happened in human history—as his explicit topic. Further, he is the first to appreciate it not just as an intellectual phenomenon, namely Enlightenment. He was the first to conceptualize the economic, political, and social transformations as all of a piece with the intellectual ones. He understood the transition from traditional to modern life, as Foucault puts it, as including “elements of social transformation, types of political institution, forms of

⁹² “What is Enlightenment?”.

knowledge, projects of rationalization of knowledge and practices, technological mutations.”⁹³

For Hegel, reason shows itself as having the form of a vast metanarrative, rationally reconstructing the emergence of modernity in all its multifarious aspects. That narrative is progressive and triumphalist. It is the emergence of reason as sovereign both in individual subjective self-consciousnesses and in the social institutions that they shape and that shape them. It is also, and essentially, as Hegel says, the “history of the progress of the consciousness of freedom.” Here two strands of the Enlightenment come together: faith in the sovereignty of reason, and the narrative of the emerging *self-conscious realization* of that sovereignty, which is the emancipatory power of reason.⁹⁴ Freedom takes concrete form only in the practical (including institutional) appreciation of the *rational* nature of genuine authority: the idea that reasons alone are normatively authoritative. This is reason’s *disenchantment* of the subordination model of authority, in favor of the model of autonomy as consisting in acting for *reasons*.

This intoxicating identification of freedom and the authority of critical reason is the beating heart of German Idealism. In it, ideas that in retrospect could be seen to have been all along implicit in Enlightenment rationalism come to fully explicit theoretical self-consciousness. It is in just such a context, Hegel thinks, that countercurrents of thought become visible as also

⁹³ "What is Enlightenment ?" ("Qu'est-ce que les Lumières ?"), in Rabinow (P.), éd., *The Foucault Reader*, New York, Pantheon Books, 1984, pp. 32-50.

⁹⁴ “But even though man, who in himself is rational, does not at first seem to have got further on since he became rational for himself—what is implicit having merely retained itself—the difference is quite enormous: no new content has been produced, and yet this form of being for self makes all the difference. The whole variation in the development of the world in history is founded on this difference. This alone explains how since all mankind is naturally rational, and freedom is the hypothesis on which this reason rests, slavery yet has been, and in part still is, maintained by many peoples, and men have remained contented under it.” Hegel’s Introduction to his *Lectures on the History of Philosophy* E.S Haldane and Frances H. Simson, trans. [Humanities Press, 1974], p. 21.

having been all along implicit in that same tradition. In this case, a crucial trajectory of nineteenth century thought expresses the revenge of Enlightenment *naturalism* on Enlightenment *rationalism*. The form that revenge took is *genealogy*. Genealogies directly challenge the very idea of the normative force of the better reason, which lies at the core of the Enlightenment rationalist successor to the traditional subordination model of authority.

The principal practitioners of the genre I am calling “genealogy” were the great unmaskers of the nineteenth century: Marx, Nietzsche, and Freud. (Closer to our own time, we might add Foucault.) *What* they unmasked were the pretensions of reason. Kant had rigorously enforced the distinction between *reasons* and *causes*, criticizing Locke for producing “a mere physiology of the understanding,” rather than a proper epistemology, by running together issues of *justification* and *causation*. We must separate, he insisted, the *quid juris*, the question of right, from the *quid facti*, the question of fact. The first is a matter of the *evidence* for our beliefs, the second of their matter-of-factual *origins*.

When the great genealogists dug down in the areas of discourse they addressed, they found *causes* underlying the *reasons*. Their enterprises can be rendered in relatively moderate terms: what they diagnosed were “systematic distortions in the structures of communication,” as Habermas puts it.⁹⁵ For Marx, the distorting causes were economic classes (functional roles with respect to the relations of production). For Nietzsche, they were expressions of the will to power. For Freud they were such things as lingering echoes of a child’s role in the Family

⁹⁵ In Jeremy J. Shapiro (trans.) *Knowledge and Human Interests* [Beacon Press, 1972].

Romance. On the moderate understanding of genealogy, those causal factors shape the reasoning of those subject to them, operating behind their backs, so that their own thoughts and actions cannot be transparent to them. This way of thinking about things leaves open the possibility of emancipatory critical discourses, which would make explicit those distorting causal factors, so breaking the hold they have on reasoners and moving them towards the ideal of rational self-transparency.

I will be concerned here, though, with a more radical challenge genealogy can be seen to make to the Enlightenment's idea of reason. For one can take it that what the genealogists dug down to is not just causes *distorting* our reasons, but causes *masquerading as* reasons. When what we fondly believe to be reasons are unmasked, all that remains is blind causal processes. Those processes have taken on the guise of reasons, but in fact yield nothing more than rationalizations. Genealogy in its most radical form seeks to dispel the *illusion* of reason.

As I shall use the term, genealogical explanations concern the relations between the *act* or state of *believing* and the *content* that is *believed*. A genealogy explains the advent of a belief, in the sense of a *believing*, an attitude, in terms of contingencies of its etiology, appealing exclusively to facts that are *not* evidence, that do *not* provide reasons or justifications, for the truth of what is *believed*. In this sense, when it occurs to the young person that he is a Baptist because his parents and everyone they know are Baptists, and that had he been born into a different community he would have with equal conviction held Muslim or Buddhist beliefs, that is a genealogical realization. As is evident already in this mundane example, the availability of a

genealogical explanation for a constellation of beliefs can have the effect of undercutting its credentials as something to which one is *rationally* entitled. The genealogy asserts counterfactual or subjunctive conditionals linking the possession of certain beliefs (attitudes of believing) to contingent events whose occurrence does *not* provide *evidence* for (the truth of) what is believed. If the believer had not had a bourgeois upbringing, were not driven by resentment, or had not had that childhood trauma, she would not have the beliefs about the justice of labor markets, Christian ethics, or conspiracy theories that she does. None of those events, upon which, the genealogist asserts, the holding of the beliefs in question are counterfactually dependent, provide *evidence* for what is believed.

For the particular vocabularies they address, all of Marx, Nietzsche, and Freud offer natural histories of the advent of beliefs (believings) couched in those vocabularies—ones to which the rational credentials of the beliefs (what is believed) are irrelevant. Natural, causal processes of belief-formation are put in place of rational ones. To him who looks on the world reductively, the world looks reductively back. This movement of thought, too, comes with its native metanarrative of progress in understanding. Replacing *theological* necessity with *rational* necessity as the fundamental explanatory category is *disenchantment* of the world **by** reason. Replacing *rational* necessity with *natural* necessity is *disillusionment* **with** reason. From the genealogical point of view, the Enlightenment apotheosis of reason just substituted one ultimately supernatural self-delusion for another.

The Enlightenment was right to be impressed by the rise of the new science, to see it as requiring a thorough-going transformation of our understanding of our relations to our world. But from the genealogical point of view, it was insufficiently radical. It naturalized, and so disenchanting the world—but it did not disenchant *us*. The Enlightenment conception of the individual knowers and agents who brought about and were in turn transformed by the convulsions of modernity retains a spark of divinity, in the form of the faculty of reason. The genealogical movement of thought teaches by contrast that the *subjects*, and their relations to the objects they know about and act on, no less than those *objects* themselves, must be thoroughly naturalized.

But what about the normative “force of the better reason”? Is it, too, just an illusion arising from the play of natural forces? Or can it somehow be understood in terms of them? Can we really understand the natural science that is the source of genealogies of our believings itself entirely in naturalistic terms? Must we? In its most radical form, the genealogical thought is that *if* we can understand the etiology of our believings (and preferring, intending, and so on) in terms of *causes* that do not provide *reasons* for them, *then* talk of reasons is shown to be out of place: not only superfluous, but actively misleading. The metanarrative of genealogy as unmasking illusions of reason depends on the disjunction “causes or reasons” being *exclusive*, its forcing a choice on us. Genealogy turns Kant’s distinction back on itself. It becomes a snake, poisoning itself by biting its own tail.

II. Global Reductive Genealogies and Semantic Naiveté

Marx and Freud offer *local* genealogies. That is, they offer genealogical analyses only of a specific range of discursive practices, the use of only some vocabularies—the vocabulary of political economy (with ripples through the cultural “superstructure”, to be sure), or the vocabulary one uses to explicate and make intelligible one’s psychological motivations. Though Nietzsche’s most detailed stories are of this local kind, he also points the way to the possibility of a more global genealogical lesson: that a suitably thoroughgoing reductive naturalism might undercut the rational credentials not just of some parochial region of our belief, but of the whole realm. The very idea of reason as efficacious in our lives would be called in question by globalizing the genealogical enterprise to extend it to *all* discourse. In this form, genealogy would be

The little rift within the lute

That bye and bye shall make the music mute,

And, ever widening, slowly silence all

as Tennyson has it. This global genealogical challenge would come to be expressed explicitly in various forms in the twentieth century, but the neo-Kantian Windelband could already find it implicit in the aspirations of his nineteenth century historicist opponents.⁹⁶ It is this broader idea that I want to consider.

⁹⁶ This story is recounted by Frederick Beiser in his chapter on Windelband in *The German Historicist Tradition* [Oxford University Press, 2012].]

Globalized genealogical arguments take a common form. They present causal etiologies of states and events of believing, thought of as episodes in the natural world, as rendering superfluous and irrelevant appeal to reasons that normatively entitle believers to the contents believed. The thought is that *all* the explanatory work can be done by *causes*, with no work left to be done by *reasons*. As a second, subsidiary, task, one then explains the motives for which and the structures by which believers and theorists conceal, from themselves and from others, the underlying causal processes of belief acquisition under an obscuring veil of what then show up as mere rationalizations. All the great genealogists of the nineteenth century particularly relished offering such metagenealogies. That is how they unmask our conception of ourselves as *rational* animals as nothing more than an illusion that puffs up and comforts animals with the sort of natural needs and interests we have. Our need for that swaddling illusion reveals us to be in essence not, as we are pleased to think, autonomous *rational* animals, but merely needy, insecure, *rationalizing* animals.

At this level of generality, the genealogical challenge to reason has the form of a naturalistic reductionism about the essentially normative “force of the better reason.” I think it is illuminating to compare this global challenge with the more focused version Gilbert Harman addresses to specifically *moral* normativity.⁹⁷ He argues that the best explanation, indeed, a complete explanation, of why people have the moral normative *attitudes* they do—why they treat some acts as morally right and others as morally wrong—need appeal only to other normative

⁹⁷ Gilbert Harman *The Nature of Morality: An Introduction to Ethics* [Oxford University Press, 1977].

attitudes (of their own, and of others). It need *not* appeal to norms or values in the form of facts about what *is* actually morally right or wrong. He contrasts this situation with that concerning our attitudes towards electrons, the best explanation of which, he takes it, must include reference to facts about electrons and our interactions with them. He concludes that we do not in the end have reason to believe in the existence of moral norms or values, as we do for the existence of electrons.

A global version of this argument, addressed to the norms of reason rather than of morality, would contend that a complete explanation for people taking or treating some claims as reasons for others need appeal only to their attitudes of *taking* or *treating* some claims as reasons for others, not to any facts about what really *is* a reason for what. Propositional attitudes, paradigmatically beliefs, would be treated just as features of the natural history of creatures like us, and hence as explicable entirely in terms of other such features, in this case, further attitudes. For our purposes, it won't matter much just what form the envisaged explanation takes, so long as it is naturalistic and does not appeal to irreducible normative facts about rational, evidential, or justificatory relations among the contents of the attitudes.

I think there is a structural defect that afflicts global reductive genealogical stories of this kind. They depend on what I will call "semantic naiveté." Semantic naiveté consists in taking for granted the conceptual *contents* of the attitudes whose rational relations to one another one wants to dissolve genealogically. If the attitudes in question are not thought of as propositionally contentful, then the issue of *rational* normative relations between them, of some of them as

providing good reasons entitling or committing one to others, does not even arise—as it does not for whirlpools, thunderstorms, supernovae, and other natural occurrences into whose causal antecedents we might inquire. The question I see as posing a counterchallenge to genealogical challenges to the very idea of reason is whether and how one is to understand the *contentfulness* of beliefs apart from their situation in a normative space of reasons.

The overall point is that *epistemological* claims, including skeptical ones, have *semantic* presuppositions. I am going to argue that the soft underbelly of genealogical skepticism about reason is its implicit commitment to a naïve semantics. When we look at things more closely, we will see that the underlying issue concerns the relations between contingencies governing attitudes—what applications of concepts are taken or treated as correct according to the prevailing reasons—and norms to which those attitudes are subject. Observations about the former provide the basis for the genealogical challenge to the intelligibility of the latter. The particular form of semantic naiveté I will identify as crucial to this debate turns out to be an assumption about the relations between semantic attitudes and semantic norms that is common both to Enlightenment rationalism and to the genealogical challenges to it. The thinker who diagnoses this shared presupposition, contests it, and offers a constructive alternative is Hegel—whom I will argue both anticipates, and responds creatively to, the genealogical currents of thought he inspired and in many ways made possible.

III. Discursive Norms and Attitudes and the Threat of a Norm/Fact Dualism

I have described global genealogical challenges to our understanding of ourselves as rational both as rooted in Kant's distinction between reasons and causes and as expressing the revenge of that distinction on itself. This is of course a very crude formulation. To refine it, we need to fill in some of the Kantian background. Kant brought about a revolution in our understanding of the mind by recognizing the essentially *normative* character of the discursive. In a decisive break with the Cartesian tradition, he distinguishes judgments and intentional actions from the responses of nondiscursive creatures not *ontologically*, by their supposed involvement with an ultimately spooky kind of mental substance, but *deontologically*, as things their subjects are in a distinctive way *responsible* for. What we believe and what we do express *commitments* of ours. They are exercises of a kind of *authority* characteristic of discursive creatures. Responsibility, commitment, authority—these are all *normative statuses*. Concepts, which articulate discursive acts of judging and intentionally doing, Kant says, are *rules*. They are rules that determine *what* we have made ourselves responsible for, *what* we have committed ourselves to, *what* we have invested our authority in. Appreciating the rulishness of the mind is Kant's normative turn.

Practically, what we are responsible for and committed to *doing* in investing our authority in how things are or are to be, Kant thinks, is having *reasons* for those commitments. What concepts are rules for *doing* is reasoning. It is the concepts articulating the contents of our judgments and intentions that determine what count as reasons for and against thinking or acting that way: what would *entitle* us to do so or *justify* us in taking on commitments with those

conceptual contents. As discursive creatures, we live and move and have our being in a *normative space of reasons*.

After Descartes, the challenge was to find a place for mental stuff in a natural world of physical stuff. After Kant, the challenge became finding a place for norms in a natural world of facts. Descartes has been roundly criticized for his dualism of minds and bodies. The danger is that the result of Kant's revolutionary insight into the normativity of intentionality would be to replace that dualism with a dualism of norm and nature. (Indeed, that is arguably the rock on which nineteenth-century German neo-Kantianism foundered.⁹⁸) I take it that a *distinction* becomes a *dualism* when it is drawn in terms that make the relations between the distinguished items unintelligible. I will argue that the collision between the possibility of global genealogies and understanding ourselves as rational depends on a set of assumptions (which can be gathered together under the rubric "semantic naiveté") that would turn Kant's distinction into a dualism--but that those assumptions are optional, and indeed incorrect. I will argue further that Hegel—intense and insightful reader of Kant that he was—already understood all this and offered a constructive alternative that can provide a way forward for us in thinking about these issues today.

Kant's normative turn expressed an insight in discursive *pragmatics*: our understanding of what we are *doing* in judging and acting intentionally. He also moved beyond the Cartesian

⁹⁸ Fred Beiser argues for this claim in "Normativity in Neo-Kantianism: Its Rise and Fall" *International Journal of Philosophical Studies*, Vol. 17, No. 1 (2009): 9-27.

tradition he inherited in seeing that its characteristic *epistemological* concerns raised a more fundamental *semantic* question. His idea here was that if we properly understood what it is for our thoughts to be representations in the sense of so much as *purporting* to represent something (to “have objective validity,” in his jargon, the successor concept to Descartes’ “*tanquam rem*”—being as if of things), the epistemological skeptical question of what reason we have to think that they ever *correctly* represent something would be revealed on semantic grounds to be ill-posed (his “Refutation of Idealism”). Hegel saw, however, that as penetrating as these archaeological semantic excavations were, Kant failed to appreciate and address a crucial *semantic* question raised by his original normative *pragmatic* idea. Kant correctly saw judging and acting intentionally as exercises of authority that come with correlative responsibilities: commitments to having reasons for and acknowledging consequences of those undertakings. He understood concepts as functions of judgment, in the sense of rules that determine what would count as a reason for applying those concepts in judgment, and what the further consequences of doing so are. In a strict sense, all Kantian rational creatures can do is apply concepts, in judging and acting. So those discursive activities presuppose the availability of the concepts they deploy. But that presupposition raises in turn the question faced by Kant’s rationalist hero Leibniz: where do those concepts come from? (The two thinkers are together in rejecting responses in terms of abstraction, which are characteristic of empiricism.) Once the discursive enterprise is up and running, new concepts can be formed downstream from applications of old ones (e.g. by “judgments of reflection”). But what is the origin of the concepts that make empirical and practical discursive activity possible in the first place?

Hegel reads Kant as having a two-stage story: transcendental activity is the source of the conceptual norms that then govern empirical discursive activity. The empirical self accordingly always already finds itself with a stock of determinate concepts. The (transcendental) processes by which discursive norms are *instituted* are sharply distinguished from the (empirical) processes in which those discursive norms are *applied*. In the twentieth century, Rudolf Carnap (in this regard, as in others, showing the effects of his neo-Kantian antecedents) provides an index example of this Kantian two-stage semantic-epistemic explanatory strategy. In his version, the two stages correspond to beginning by fixing *meanings* and only then fixing *beliefs*. The first, semantic, stage is selecting a *language*. The second, epistemic, stage is selecting a *theory*: a set of sentences, couched in that language, that are taken to be true. His student Quine objected to Carnap that while this two-stage procedure makes perfect sense for formal or *artificial* languages, it makes no sense for *natural* languages. All speakers do is *use* the language—Kant would say, to make judgments. That use must somehow determine both what their expressions mean and which sentences they take to be true. In the vocabulary I used to talk about Kant, the use of language to express judgments must be understood as effecting *both* the *institution* of conceptual norms *and* their *application*.

Two-stage stories about the division of labor between semantics and epistemology—that is, about the relations between conceptual contents and their application in judgment—are committed to *semantic purity*. This is the view that the *contents* concepts possess are not at all affected by the *use* of those concepts in making judgments: believing a particular subset of the universe of believables. That is the point of having a first, semantogenic stage at which contents are determined, conceptual norms instituted, before second stage comprising the application of

the concepts in taking things to be thus-and-so—to be as represented by some already contentful representings, and not others. Commitment to semantic purity is commitment to the possibility of pursuing semantics independently of commitment to how things actually are. The thought is that *epistemic* commitments are not to contaminate *semantic* ones. Semantic commitments are necessary conditions only for the *expression* of epistemic ones. On this picture, two independent elements combine to make epistemic commitments (true claims): semantic commitments (picking a language, concepts), and how the world is. The second element is irrelevant to the first. (In recent years, we have come to doubt this picture on the basis of arguments, by Putnam, Burge, and Davidson, among others, for what has come to be called "semantic externalism".)

Semantic purity is not an unintelligible idea. It makes sense in the context of stipulating associations of semantic interpretants with linguistic expressions for an artificial language, by a theorist working in a semantically more powerful metalanguage. Semantic *naïveté* results when one believes that semantic *purity* is intelligible for an *autonomous* intentional stratum: for natural languages, or for thought in general. Quine objects to the semantic naïveté of commitment to the possibility of pure semantics—and in this regard makes common cause with the later Wittgenstein. Both thinkers take it that all there is to confer *content* on our expressions is the way those expressions are *used*: paradigmatically, in making claims and forming beliefs, that is, in committing ourselves to how things are. Two-stage theories about the division of semantic and epistemological labor for natural languages and the thoughts they express, they think, are bound to invoke semantic stories about the first stage that make the notion of conceptual content ultimately magical. They are committed to semantic purity. So, when applied to natural languages and thought, they are semantically naïve. This is exactly Hegel's complaint about

Kant: he was uncharacteristically, but culpably, uncritical about the source and nature of determinate conceptual contents. In this regard, Hegel is to Kant as Quine is to Carnap. And like Quine and Wittgenstein, Hegel offers an ultimately *pragmatist* account of how using a natural language can be intelligible as both instituting and applying conceptual norms.

This line of thought bears directly on the issue we are considering. For global genealogical reductive explaining away of *norms* in favor of *attitudes* presumes that it is intelligible for the *contents* of propositional attitudes to stay in place after normative reason-relations among their judgeable contents are relinquished. Otherwise what is being explained genealogically can no longer be understood as *believings*—as attitudes of taking things to be (representing them as) thus-and-so. If our attitudes were not genuinely conceptually contentful, then we would not even be *purporting* to represent things as being thus-and-so; things would not even *seem* to us to be thus-and-so. If disillusionment about the reality of norms of reasoning entails *semantic nihilism*, then it is self-defeating: the genealogists claims would entail that her own claims are senseless. (Here we might compare the argument that if we were brains-in-vats, then nothing could give our thoughts the content either that we are brains-in-vats, or that we are *not* brains-in-vats—or indeed, any propositional contents at all.) I think in fact there is a good inferentialist argument to the conclusion that no account according to which discursive practitioners never actually are rationally entitled to their claims, justified in holding them, have real evidence for them (all normative statuses) can ultimately make sense of their normative attitudes having semantic content. But I am not going to rehearse such an argument here.

The point I want to make is that taking the contents of propositional attitudes in general to be independent of the government of those attitudes by norms concerning what is genuinely a reason for what presupposes a semantically naïve two-stage account of the division of semantic and epistemic labor. For it requires that the contents of propositional attitudes have already somehow been fixed in advance and independently of the rough-and-tumble of assessing evidence and deciding what to believe. The semantic challenge for the globalized Harmanian genealogist is accordingly to say how we are to understand the *contents* of the *attitudes* in favor of which genuine norms have been eliminated. The corresponding challenge for a one-stage account is to explain how institution of genuine conceptual norms is compatible with the possibility of genealogical explanation of acts of applying such norms. Hegel understands this challenge, and offers an intricate and sophisticated story about the relations between the institution and the application of conceptual norms, including the relations between discursive normative statuses and discursive normative attitudes, that is aimed precisely at meeting that challenge. In the rest of this talk, I want to present the outlines of that story, as I understand it.

IV. From *Verstand* to *Vernunft*

One way into Hegel's constructive alternative to the semantic naivete of two-stage theories of the division of semantic and epistemic labor is through his conception of the *determinateness* of conceptual norms. What semantic purity claims conceptual contents are pure *of* is contamination by the *epistemic*, that is, by *knowledge* claims, *judgments* as to how things *actually* are. The *semantics* of concepts (universals) is supposed not to depend at all upon

epistemic commitments, that is, on judgments. (A prominent avatar of commitment to semantic purity in our own time is Jerry Fodor, who insists that running together semantic and epistemic issues—as he sees Davidson and Dummett, as doing, for instance—is the methodological Great Bad that must be overcome if semantics is to move forward.) Hegel’s slogan for the conceptual sea-change he sees as necessary (and sufficient) for appreciating the interdependence of semantic and epistemic commitments is that we must move from understanding the conceptual in terms of static categories of *Verstand* to understanding it in terms of dynamic categories of *Vernunft* (adapting Kant’s terminology to his own uses).

Kantian concepts are determinate in the *Verstand* sense in that the rational relations of consequence and incompatibility between concepts (universals), which identify and individuate them, are taken to be fully settled in advance of any application of those universals to particulars in judgment. Kant envisaged an asymmetric structure of capacities, in which a faculty of *spontaneity* (activity) is the source of universals, which are applied to the particulars supplied by a faculty of *receptivity* (passivity). In developing his successor *Vernunft* conception, Hegel takes over from Kant his insight into the normative character of concept-use, and radicalizes it by construing the relations between universals and particulars itself in normative terms of authority and responsibility (his “independence” and “dependence”). Hegel takes his cue from the fact that, transposed into the normative key, the relations of authority and responsibility between universals and particulars are reciprocal and symmetric. Kant’s system masks that underlying symmetry by an artificial, asymmetric division of semantic and epistemic labor. Spontaneous exercises of the *semantic* authority of the understanding (*Verstand*) over universals are independent of and prior to exercises of the *epistemic* authority of particulars (in the form of the

manifold of intuition delivered by receptivity) which determine the *correctness* of applications of universals to those particulars in judgment. This overarching asymmetric structure is a manifestation of Kant's understanding of the freedom of reason in terms of *autonomy* (pure independence).

According to Hegel's symmetric normative construal of the relations of authority and responsibility between universals and particulars, the application of one concept (universal) obliges one to apply others to that particular (according to relations of rational *consequence* that articulate (partially determine) the content of the concept=universal) and precludes one from being entitled to apply others to that particular (according to relations of rational *incompatibility* that articulate (partially determine) the content of the concept=universal). This is the authority of universals over particulars, the responsibility of particulars to universals. (Kant understood this as a semantic matter.) There is a corresponding relation of authority of particulars over universals. For it can happen that one applies a concept (universal) to a particular and the particular does not cooperate in also exhibiting the universals that are its consequence, or in also exhibiting universals that are incompatible with the original one. This Hegel construes as the particular exercising authority over the universal: telling it, as it were, that it cannot have the consequence-and-incompatibility relations that it originally came with, that is, that a different universal is required. (Kant understood this as an epistemic matter.) For Hegel, none of these reciprocal relations of authority and responsibility between universals and particulars should be understood as purely semantic nor as purely epistemic. The clean division of semantic and epistemic labor is an artifact of semantically naïve two-stage accounts. Our judgments shape our concepts no less than our concepts shape our judgments.

Hegel understands *determinateness* (Bestimmtheit) in terms of what he calls “*individuality*” (Einzelheit). Individuality, in turn, is a matter of the characterization of a *particular* by a *universal*, which is something that has the form of a *fact* or a *judgment* (in the sense of judgeable content, which, when true, is a fact). As Kant emphasized, concepts shape and articulate judgments. Hegel adds the idea that judgment is the process by which concepts are determined. The essence of Hegel’s *Vernunft* conception is an account of the structure of the dynamic *process* in which the whole constellation of concepts-and-judgments (what Hegel calls “the Concept”) develops by the exercise of the reciprocal authority of universals over particulars and particulars over universals. Judging, the application of universals to particulars, is the development of *individuals*: at once the semantic shaping and determining of universals and the epistemic discovery of which universals apply to particulars.

Kant’s pure independence model of semantic authority as untrammelled by corresponding responsibility leaves it unclear what room there remains for epistemic constraint. Why cannot the boundaries (implications and incompatibilities) of the universal that has been applied simply be redrawn to accommodate any looming recalcitrance (in the mode of the preadolescent, who, falling off a bicycle, says “I meant to do that”)? More deeply, what counts as *changing* the content of the concept (universal)? What holds fixed, in advance, the commitments one undertakes by applying it, if its content is wholly up to the “spontaneous” activity of the subject? The Kantian division of semantic and epistemic labor seems unable to exclude the possibility that “whatever seems right to me is right”—in which case the issue of correctness does not get a

grip (as Wittgenstein puts the point). There is nothing in the Kantian picture to confer *determinate* contents on concepts, nor to hold them in place *as* they determinate.

What is needed, Hegel thinks, is to replace Kant's *individualistic* model (driven by his understanding of freedom as autonomy) with a *social* one. What Kant tried to accomplish within the boundaries of a single knowing subject by the division of semantic and epistemic labor should rather be done by a genuinely *social* division of labor. Concepts for Hegel are not to be found between the ears of individual knowers, but in the public *language* they speak.⁹⁹ (As Hilary Putnam would later put the point, "Meanings ain't in the head.") This transposition of the issue into a social, linguistic key makes it clear how in judging, whose paradigm now becomes asserting, I can bind myself by norms provided by the concepts I apply to particulars. It is wholly up to me whether I assert that the coin is copper—rather than manganese, say. But it then *not* up to me what else I have committed myself to by claiming that, and what would entitle me to that commitment. The metallurgical experts my community charges with the care and feeding of the concept copper will *hold* me responsible for having committed myself to the coin's melting at 1084° C., and to have precluded myself from claiming that it is an electrical insulator. Whether I *know* about these implications is neither here nor there. They are features of the move I have made in my language-game. It is my participation in that game that permits me also to *think*, quietly to myself, *that* the coin is copper—a thought that inherits its shared content from claimables whose sense the community fixes.

⁹⁹ "Language," Hegel says "is the Dasein of Geist" [the *Phenomenology of Spirit*, [A.W. Miller, (trans.), Oxford University Press] paragraph 652], that is, it is the actual existence of the normative, and hence of the discursive.

On this model, the *authority* of an individual speaker (what Kant construed as *autonomy* and Hegel as a moment of *independence*) is balanced by a reciprocal *responsibility* (a moment of *dependence*). And the content I have freely committed myself *to* and made myself responsible *for* is held in place as *determinate* by my fellow speakers, whom I have authorized to hold me responsible *for* it. (That is *their* “moment of independence”.) What I am responsible for is what I *said*, not what I might later claim to have *meant*. (The sense in which I can privately mean something other than what I say is intelligible in principle only against the semantic background of what is publically sayable.)

What Heidegger called “the dignity and spiritual greatness of German Idealism” is founded on Kant’s reconstrual of self-conscious *selfhood* as consisting in *freedom*, in the sense of the *authority to commit* oneself determinately, the capacity to *bind* oneself by conceptual *norms*, norms that are *rational* in the sense that they articulate what is a *reason* for a judgment or an action with that content. Hegel sees that self-consciousness in this normative sense is an essentially *social* achievement. The authority to *make* oneself responsible (for what one thinks or what one does) makes sense only in a context in which one can be *held* responsible. That requires *two* loci of authority and responsibility. Normative statuses, such as authority and responsibility, and the selves that are the subjects of such statuses, Hegel teaches, are instituted by *reciprocal recognition*: that is, by individuals practically taking or treating one another *as* authoritative and (so) responsible. Those I recognize, in this normative sense of authorizing to hold me responsible, form a recognitive community. In telling language, Hegel says that self-

conscious *individual* selves (normative subjects) are instituted only when *particular* organisms come to stand in recognitive relations to one another (a matter of their practical normative attitudes), and so to be characterized by the *universal* that is the recognitive community.

As I read him, at the center of Hegel's thought is the idea that there is a structural—he says “*logical*”—relation between particularity, universality, and the individuality that consists in a particular being characterized by a universal, that has as species all of

- the *semantic* relations between particulars in the sense of singular terms and the universals in the sense of predicates that characterize them in *judgments*,
- the *ontological* relations between particulars in the sense of objects and the universals in the sense of properties that characterize them in *facts*, and
- the *social-normative* relations between particulars in the sense of human organisms and the universals in the sense of their recognitive communities, in virtue of which they become self-conscious individual selves in virtue of being members of those normative communities.

Accordingly, I take it that the cardinal standard for assessing a reading of Hegel is how it understands and assesses this assimilation.¹⁰⁰ In my reading, the social recognitive species is the key one, in terms of which the other two are to be understood. This is controversial, and I will not here expound either my understanding or my assessment of the value of this central move of Hegel's.

¹⁰⁰ Given the centrality I assign to this move of Hegel's, it is surprising to me how few interpreters so much as discuss it. Clearly someone is missing something crucial. A sensible reader will accordingly take suitably seriously the possibility that it is me. I do.

V. The Normative Structure of Tradition and Reason's March through History

Besides developing Kant's normative insight along the *social* dimension, Hegel develops it along a *historical* dimension. What binds them together is Hegel's idea that determinateness on the side of the *content* of conceptual norms (the topic of *semantics*) is intelligible in principle only in the context of a thoroughgoing *reciprocity* of authority and responsibility on the side of the practical *force* or *significance* of those norms (the topic of *pragmatics*, in a broad sense). The metaconcept of *Vernunft* is a view about the *process of determining* conceptual contents, and the kind of determinateness that results. This process has the normative structure distinctive of a *tradition*. Hegel's two-dimensional approach could be denominated "Confucian", insofar as his social/historical dimensions line up with the fraternal/filial synchronic and diachronic familial dimensions on which that tradition focuses its account of determinate obligation and normative status.

Understanding genealogical analyses as undercutting the claims of reason (the rational bindingness of conceptual norms) depends on assessing the rationality of discursive practice solely on the basis of the extent to which applications of concepts, whose contents are construed as always already fully determinate, are responsive exclusively to evidential concerns.

Responsiveness of concept-application to any factors that are contingent relative to the conceptual norms already in force—the phenomenon genealogical diagnoses highlight—is accordingly identified as irrationality. But the idea that assessments of rationality are appropriately addressed only to the application of already fully determinate concepts is the product of a blinkered semantic naiveté. It ignores the fact that the very same discursive practice that is from one point of view the *application* of conceptual norms is from another point of view the *institution* of those norms and the *determination* of their contents. Only when discursive practice is viewed whole does its rationality emerge. If the semantogenic process by which conceptual contents are determined and developed is ignored, the distinctive way in which reason informs and infuses discursive practice remains invisible.

For Hegel, the principal task of reason—in his preferred sense of *Vernunft*, rather than *Verstand*—is, as he says, to “give contingency the form of necessity.” Following Kant, by ‘necessary’ he means “according to a rule.” That is, reason’s job is to put the sort of material contingencies the genealogist points out into *normative* shape. From Hegel’s point of view, then, far from *undercutting* reason, the possibility of genealogical explanation just underlines the *need* for this particular function of reason, and the crucial job it does.

How can we understand the process whereby concepts acquire and develop their determinate content as putting contingencies of their application into normative shape? Hegel’s idea is that a distinctive kind of *retrospective rational reconstruction* of prior applications of a concept is necessary and sufficient to exhibit those applications as conferring a determinate content on the

concept. (This is what in the final chapter of the *Phenomenology* he calls ‘Erinnerung.’) One brings order to the motley welter that is the discursive practice one inherits by discriminating within it a privileged trajectory that is *expressively progressive*, in the sense of making gradually *explicit* norms that then show up as having been all along *implicit*. Doing that is turning a *past* into a *history*.

The best model I know of the kind of rational activity that determines conceptual contents by making or finding the right kind of history for them is the *jurisprudential* one institutionalized and codified in *case law*. Its purest, paradigmatic form takes place in what in the Anglo-American legal world is called “common law.” For in that realm, by contrast to statute law, judges are not guided in their decisions as to whether to apply or to withhold application of a concept (such as strict liability) by explicit statutes, propounded and given the force of law by legislatures—statutes that *say* what is and is not licit according to the norm they institute. In lieu of norms *explicit* as such *principles*, judges at common law must decide cases with novel facts on the basis only of norms they discern as *implicit* in the *tradition* of already-decided cases.

The governing authority to which common law judges are responsible is provided by *precedent*. The judge’s job is not only to decide the present case, but also to provide a *rationale* for the decision, by providing a distinctive kind of narrative justifying it as correct. Such a narrative selects some prior decisions as precedential, in the sense of being not only relevant and correct, but as having revealed some hitherto hidden aspect or contour of the norm developing in the tradition defined by those precedents. The legal concepts and the principles explicating them

that are given expression in rationales for deciding novel cases are often characterized as “judge-made law.” This description is apt, because there is nothing more to give content to this kind of law than the decisions judges have rendered and the retrospective rational reconstructions of traditions defined by precedent that the judges offer to justify those decisions.

Rational, rationalizing processes of this sort both are *responsible to* the contents of the conceptual norms they apply, and they exercise *authority over* the development of those conceptual contents. They are processes of *determining* conceptual contents both in the sense of *finding out* what they are, manifested in the essentially *retrospective* rationales judges supply for their decisions, and in the sense of *making* those contents what they are, manifested in the essentially *prospective* shifting and sharpening of the norms each new application and interpretation proposes. These hermeneutic practices give contingency the normative form of necessity, and by incorporating those contingencies infuse determinate content into the developing norms. It is of the essence of the kind of rationality distinctive of this sort of concept-determining process to be articulated by these complementary perspectives: *retrospective determining-as-finding* and *prospective determining-as-making*, *responsibility to* the tradition one inherits and *authority over* the tradition one bequeaths. Looking *backward* along the privileged trajectory of precedents selected by the narrative rationalizing any particular decision, one sees only unbroken expressive progress: the gradual emergence into the explicit light of day of a governing norm that appears as having been all along implicit in earlier decisions. Looking *forward* at how legal concepts and principles evolve by being applied in concrete cases, the *discontinuities* between these narratives shows up, as sequential judges revise their predecessors’ judgments as to which earlier applications should be treated as precedential,

and how. T.S. Eliot describes this aspect of Hegelian *Vernunft* at work in a different corner of the culture:

What happens when a new work of art is created is something that happens simultaneously to all the works of art which preceded it. The existing monuments form an ideal order among themselves, which is modified by the introduction of the new (the really new) work of art among them. The existing order is complete before the new work arrives; for order to persist after the supervention of novelty, the whole existing order must be, if ever so slightly, altered; and so the relations, proportions, values of each work of art toward the whole are readjusted; and this is conformity between the old and the new.¹⁰¹

The normative structure of tradition epitomized in the development of the concepts of common law is a diachronic version of reciprocal recognition. Each judge recognizes the authority of earlier judges and petitions for recognition from those to come. All exercises of discursive interpretive authority come with commensurate responsibilities, administered by those whom one petitions for recognition of that authority—who are in turn responsible for their assessments, which can be overturned on the authority of still future judges. I think this is a model (stylized, to be sure) for how Hegel thinks the contents of *all* ground-level empirical and practical concepts develop and are determined—how conceptual norms are instituted as they are being applied.

¹⁰¹ In “Tradition and the Individual Talent” [ref.]

VI. From Semantic Naiveté to Hermeneutic Magnanimity

Considering genealogical counterfactuals about what the norms would have been had various non-evidential factors differed reveals a judicial process shot through with contingencies—as, for instance, where the order in which two cases happened to be adjudicated evidently affects the content of the law that results. The normatively contingent character of any particular decision to apply or not to apply a particular concept is manifested in the fact that one always *can* explain any particular decision *genealogically*—in terms of “what the judge had for breakfast,” in the derisive slogan of jurisprudential theory. That is to explain it in terms that do not appeal to the *content* of the norm whose applicability is in question—to explain it instead for instance in terms of the intellectual fashions or public passions of the day, or by features of the judge’s training, temperament, or political convictions. But to conclude that the possibility of such an explanation means that no norm is therefore instituted, that that the norms discerned as implicit in the tradition inherited cannot *rationally* justify one decision rather than another in a novel case, is to insist stubbornly and one-sidedly on occupying only one of the perspectives that are in fact two sides of one coin—as Hegel insists, and jurisprudential practice demonstrates. It is precisely to refuse to see *Vernunft* whole. It is to embrace the semantic naiveté that ignores the essential role rationally incorporating those contingencies plays in conferring determinate content on (determining the content of) always evolving conceptual norms.

Hegel points to the generality of the lesson he wants us to learn from his *Vernunft* model of the practice of reason in a remarkable passage epitomized in an aphorism expressing his twist on a slogan of his day:

“No man is a hero to his valet;

not, however, because the man is not a hero, but because the valet—is a valet....”¹⁰²

The passage continues, explaining that the reason is that the valet’s

...dealings are with the man, not as a hero, but as one who eats, drinks, and wears clothes, in general, with his individual wants and fancies. Thus, for the judging consciousness, there is no action in which it could not oppose to the universal aspect of the action, the personal aspect of the individuality, and play the part of the moral valet towards the agent.

What Hegel calls the “universal aspect of the action” is its normative dimension. The hero is a hero insofar as he acts according to the norms that articulate his duty. The valet views what the hero does genealogically, in resolutely naturalistic, nonnormative, reductive terms, and so

...explains [the action] as resulting...from selfish motives. Just as every action is capable of being looked at from the point of view of conformity to duty, so too can it be considered from the point of view of the particularity [of the doer].... If the action is accompanied by fame, then it knows this inner aspect to be a desire for fame.....[T]he inner aspect is judged to be an urge to secure his own

¹⁰² *PG* [665].

happiness, even though this were to consist merely in an inner moral conceit, in the enjoyment of being conscious of his own superiority and in the foretaste of a hope of future happiness. No action can escape such judgement, for duty for duty's sake, this pure purpose, is an unreality; it becomes a reality in the deed of an individuality, and the action is thereby charged with the aspect of particularity.

Here Hegel, writing in 1806, before the advent of the great unmaskers of the dawning nineteenth century, acknowledges that *every* application of a norm is in principle liable to a naturalistic, genealogical explanation. It *can* be seen, indeed seen *correctly* (as far as that vision reaches), from the point of view of its particularity, its normative contingency. But that valet's-eye genealogical view is one-sided, it fails to see the whole of the doing. For the valet fails to see that a norm can *also* be active, that the particular contingent motives he sees (what the hero had for breakfast) can be given the form of normative necessity, can be incorporated in a narrative that exhibits them as in conformity to duty, as correctly performed according to the governing norms.

Hegel's whole metaphysical, metaconceptual apparatus is couched in the logical language of 'particularity', 'universality', and 'individuality'. Individuals are particulars as characterized by universals. Particularity and universality are to be understood as abstract aspects of concrete individuality, intelligible as determinate only in their interaction. This is not the place for me to embark on a disquisition on how I think we should understand this idiom. It suffices for present purposes to say that what stands behind Hegel's deployment of these logical categories is the story about the Janus-faced *vernünftig* process of giving contingency the form of

necessity, for which I have suggested the paradigm of the adjudication of cases at common law. The use of the logical vocabulary in characterizing the failings of the one who plays the moral valet to a hero of duty marks the extreme generality of the phenomenon Hegel is addressing, and of the conclusion he wants us to draw.

Hegel calls the genealogical valet's attitude 'Niederträchtigkeit': literally, something like a striving for the low, an impulse to debase. His term for the practical attitude of giving contingency the normative form of necessity is 'Edelmütigkeit': magnanimity. It is a form of norm-instituting recognition—the final form he discusses in the *Spirit* chapter of the *Phenomenology*—and hence a form of self-consciousness. Its retrospective recognitive aspect he calls 'forgiveness' Its prospective recognitive aspect he calls 'confession.' *What* one forgives is the normative contingencies that infect prior applications of concepts. One forgives them not wholesale, by a grand gesture, but by the hard retail work of constructing an expressively progressive historical narrative in which they play precedential roles as making explicit aspects of the developing conceptual content that are now revealed as hitherto having been implicit. The slogan of this generous hermeneutic recognitive attitude is Tennyson's:

Yet I doubt not through the ages one increasing purpose runs,
And the thoughts of men are widened with the process of the suns.¹⁰³

This is what Hegel does in his own writings—first and foremost in the *Phenomenology*, but no less in the histories he constructs for philosophy, art, and religion in his lectures.

¹⁰³ "Locksley Hall" (1835) line 137.

Concrete magnanimous hermeneutic forgiveness is *finding* such a “purpose”—that is, a norm—to which the concept-application being forgiven can be seen to contribute, widening the thoughts of man. Hegel calls this “making what happens into something *done*.” What the magnanimous interpreter confesses is the contingent inadequacy of each particular such forgiving rational reconstruction. One confesses that one is unable to find a narrative in which *every* contingency is given the normative status of a progressive precedential expression of the underlying developing conceptual norm. In confessing, one petitions one’s successors for forgiveness of that contingent failure of one’s own efforts at forgiveness. The edelmütig rational, rationalizing process in which conceptual norms are instituted by diachronic magnanimous reciprocal recognition is a structure of *trust*: trust that one’s trespasses will be forgiven as one forgives those who have trespassed before one.

So Hegel foresaw the genealogical challenge to rational normativity that would arise from a reductive naturalism and would result in a small-minded, niederträchtig valet’s hermeneutics of *suspicion*. The hermeneutics of *magnanimity* and *trust* he recommends instead is not based on fine feeling or pious sentiment. Instead he argues that the only construal on which reason and meaning are threatened by the possibility of genealogy is a narrow, one-sided conception that is mistaken because semantically naïve. In its place he puts a more capacious conception of *Vernunft* as comprising not only the norm-governed application of concepts but the process and practice by which their

content is determined. At its core is the magnanimous hermeneutics that shapes genealogical contingency into a normative, rational form. My aim here has been to sketch in the broadest outlines the insights that underlie Hegel's inspiring vision of the relations between the normative and the natural.

End