

ONE



Semantic Inferentialism and Logical Expressivism

I. Introduction

I want to introduce here 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 will 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 then 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 pick up strands other than those usually emphasized when we read 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” (Spinoza) 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*. There is, however, 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's 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 cor-

rectnesses 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 goes, 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’s 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 cannot 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 sixty 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 do not 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 does not 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 judgments, 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 “Empiricism and the Philosophy of Mind,” 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 to pursue not 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 representings 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 1890s. 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:

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.¹

Two claims have the same conceptual content if and only if 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 nonvalid sentences which are its consequences (that is, 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 1890s makes truth, rather than inference, primary in the order of explanation. Dummett says of this shift:

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.²

And again:

It remains that the representation of logic as concerned with a characteristic of sentences, truth, rather than of transitions from sentences to sentences, had highly deleterious effects both in logic and in philosophy. In philosophy it led to a concentration on logical truth and its generalization, analytic truth, as the problematic notions, rather than on the notion of a statement's being a deductive consequence of other statements, and hence to solutions involving a distinction between two supposedly utterly different kinds of truth, analytic truth and contingent truth, which would have appeared preposterous and irrelevant if the central problem had from the start been taken to be that of the character of the relation of deductive consequence.³

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 1890s 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 talk about this below under the heading of "formalism" about inference). The view propounded and attributed to Frege below is different from, and from the contemporary vantage point more surprising than, the one Dummett endorses here.

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 inferen-

tial 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 “the received dogma . . . that the inference which finds its expression in ‘It is raining, therefore the streets will be wet’ is an enthymeme.”⁴

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 involve 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:

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.”⁵

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: “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.”⁶

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 nonprivileged for 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 nontheological (or nonaesthetic) vocabulary for nontheological (nonaesthetic) 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 will 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: "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'."⁷ 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.

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 on 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 on 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 the passage quoted earlier puts it, “Whatever is needed for a correct inference is fully expressed.” Talking about this project, Frege says: “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.”⁸ The concept-script is a formal language for the explicit codification of conceptual contents. In the Preface to *Begriffsschrift*, 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.

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:

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.⁹

By contrast, Frege continues:

- i. 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 . . .¹⁰

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: “It seems to me to be easier still to extend the domain of this formula language to include geometry. We would only have to add a few signs for the intuitive relations that occur there. . . The transition to the pure theory of motion and then to mechanics and physics could follow at this point.”¹¹

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 alongside (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 one takes the fateful step from seeing logic as an attempt to codify inferences to seeing it as the search for a special kind of truth, 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: “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.”¹² 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 judgment 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 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 means 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, owing 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 principal 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 principal 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: “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.”¹³ Dummett presents his model as specifying two fundamental features of the *use* of linguistic expressions, an idea I will return to below. In what follows here, though, I apply 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.

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 that 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, at the beginning of this chapter, of Sellars's application of inferentialist ideas to the understanding of noninferential reports 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 cannot 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:

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.¹⁴

Even more clearly:

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 predicate 'valid' to an argument, and the semantic characterization of validity as giving the consequences of such an application . . . [I]f

[one] is taught in a very unimaginative way, [one] 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 [one] had missed the point of the distinction.¹⁵

Pragmatists of the classical sort, by contrast, 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 does not 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.

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

Of course, such one-sided theories do not simply ignore the aspects of content they do not treat as central. Dummett writes:

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.¹⁶

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 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 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 which it is appropriate to impose on the relation between the two sorts of rules: "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.”¹⁷ Recognition of the appropriateness of such a requirement arises from consideration of connectives with ‘inconsistent’ contents. As Prior¹⁸ pointed out, if we define a connective, which after Belnap we may call “tonk,”¹⁹ 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 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 also 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

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.²⁰

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

I have been told (by Jonathan Bennett) that 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 non-conservativeness 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 the term its 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 com-

mitments. 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 one another, 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 on 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 will not ‘tonk up’ a conceptual scheme. In a footnote, Dummett explicitly denies that conservativeness can in general be treated as a sufficient condition of harmony: “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 . . . [T]he most difficult case is probably the vexed problem of personal identity.”²¹ In another place, this remark about personal identity is laid out in more detail:

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.²²

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 focusing 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 non-logical 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 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. If we shift our concern 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 important point would be 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's 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 chapter 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 have man-

aged to say rather a lot about conceptual content in this essay, 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 chapter 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.