

*Philosophical Issues, 15, Normativity, 2005*

## **A NEW MAP OF THEORIES OF MENTAL CONTENT: CONSTITUTIVE ACCOUNTS AND NORMATIVE THEORIES<sup>1</sup>**

Mark Greenberg  
University of California, Los Angeles

### **I. Introduction**

In this paper, I propose a new way of understanding the space of possibilities in the field of mental content. The resulting map assigns separate locations to theories of content that have generally been lumped together on the more traditional map. Conversely, it clusters together some theories of content that have typically been regarded as occupying opposite poles. (Theories of content, in the relevant sense, address the question of what makes it the case that a thought, such as an act of judging that something is the case, has a particular content.<sup>2</sup>)

I make my points concrete by developing a taxonomy of theories of mental content, but the main points of the paper concern not merely how to classify, but how to understand, the theories. Also, though the paper takes theories of mental content as a case study, much of the discussion is applicable to theories of other phenomena.

To a first approximation, the difference between the traditional and the proposed taxonomies turns on whether we classify theories of content by, on the one hand, their implications for a non-redundant supervenience base for content facts (i.e., for facts about what contents thoughts have) or, on the other, by their constitutive accounts of content. By a “constitutive account,” I mean the kind of elucidation of the nature of a phenomenon that theorists have tried to give for, for example, knowledge, justice, personal identity, consciousness, convention, heat, and limit (in mathematics). An example is Locke’s view that facts about personal identity obtain in virtue of facts about psychological continuity.

The tendency to taxonomize by supervenience base is encouraged, I suggest, by a failure to keep clearly in view a distinction between constitutive and modal determination. (A contributing factor may be that the term “constitutive” is often used to draw a relatively coarse-grained distinction between the constitutive and the causal, a distinction with respect to which the modal belongs on the constitutive side.) I think that many philosophers would accept that a constitutive account cannot be captured in purely modal terms. Giving a constitutive account is not the same as specifying modally necessary and sufficient conditions. For example, being a type of polygon that can be inscribed in a circle regardless of the lengths of its sides is necessary and sufficient for being a three-sided polygon.<sup>3</sup> But a correct constitutive account of a triangle will plausibly mention the latter but not the former property.

Although such points are not unfamiliar, philosophers often try to cash constitutive claims in modal terms. A case in point is that theories of content tend to be conceptualized in terms of the theories’ implications for a supervenience base for content facts.

My thesis goes beyond the by-now somewhat familiar proposition that not all modal determinants of a phenomenon are constitutive determinants. One who has taken that point on board might nevertheless conceive of a philosophical account as an attempt to specify *constitutive* determinants of the target phenomenon that make up a non-redundant supervenience base for the phenomenon. Shoehorning a philosophical account into this form leaves out elements that are modally redundant, but may be explanatorily or ontologically significant. For example, when a constitutive account has multiple levels, the different levels will typically be modally redundant. Formulating the account as a specification of a supervenience base of constitutive determinants will therefore flatten the account into a single level.

Many of my arguments can be illustrated by considering the place of normativity in the theory of content. The new taxonomy gives a distinct niche to *normative theories of content*—theories that explain a thought’s having a certain content at least in part in terms of the obtaining of normative facts. By contrast, on a traditional map, normative theories are invisible as such because normative facts supervene on non-normative ones.

For ease of exposition, I will often talk as if there were a language of thought, using “symbol” as a generic term for bearers or vehicles of mental content. The discussion does not depend, however, on a particular view about the nature of those bearers. They might, for example, be whole propositional attitude states.

For a thinker *to have a concept* is, as I will use the term, for her to have thoughts whose contents have the concept as a constituent, or, more precisely, for her to have whatever is necessary to have such thoughts. The notion of having a concept is thus pre-theoretical. It leaves open the

question of what it is in virtue of which a symbol or thought has a particular content—the question to which theories of content propose answers. For a symbol to *express* a concept is for it to have that concept as its content.

In the next section (II), I briefly lay out a traditional way of classifying theories of mental content. The main work of the paper takes place in section III, where I develop my new taxonomy. In section IV, I conclude with a number of observations about the contrast between the traditional and new taxonomies.

## II. The Traditional Taxonomy

In this section, I give an overview of what I take to be a traditional way of categorizing theories of mental content.<sup>4</sup> Given the summary nature of the enterprise, the discussion will necessarily oversimplify. For example, there are many possibilities for hybrid theories that I do not discuss.

The most prominent classification in the traditional taxonomy is externalist versus internalist.<sup>5</sup> Pure externalist theories hold that a symbol's content is determined by its relations to objects or properties in the world outside the thinker. Pure internalist theories hold that a symbol's content is determined by the symbol's relations to other symbols in the thinker's head.

On the externalist side, the traditional taxonomy divides theories into three basic categories. Perhaps the simplest kind of externalist theory maintains that a symbol's content is determined by what in the world it covaries with. So, on a crude version of such a theory, if occurrences of a given thinker's mental symbol *S* covary with the presence of dogs, then *S* expresses the concept *dog*—i.e., has *dog* as its content. Variations of this general type of theory can hold that content is determined by causal, nomic, or counterfactual relations between symbols and worldly objects or properties. I'll refer to such theories as *indication theories* of content.<sup>6</sup> (As I use the term, it includes what are sometimes called "covariational" or "information-based" theories.)

A second, *biological* category of externalist theories holds that a symbol's content is determined by evolutionary history. An important motivation for introducing such theories is that indication theories have difficulty allowing for the possibility of systematic misrepresentation—sometimes called the *disjunction problem*. If a symbol represents whatever it covaries with, then it cannot systematically misrepresent. One response to the disjunction problem is to hold that a symbol represents whatever it covaries with under certain special conditions. Biology provides one way, though not the only one, of defining such conditions.<sup>7</sup> A different response to the disjunction problem is to hold that a symbol's content is determined by biological function, which is then explicated in terms of natural selection.<sup>8</sup>

In practice, biological theories of content tend to be categorized as externalist theories that are variations on, or close relatives of, indication theories, though writers sometimes indicate some uncertainty about this placement.<sup>9</sup>

A third kind of externalist theory is the causal-chain theory, which derives from Kripke's (1980) seminal remarks about names and natural kinds.<sup>10</sup> Most theorists recognize—as Kripke (1980, 93, 96–97) did—that an appeal to causal chains is at most a partial theory of content. Theories of various kinds incorporate causal-chain elements.

Unsurprisingly, (pure) externalist theories tend to yield contents that are quite coarse grained—that cut no more finely than reference. For theorists who believe that we need more fine-grained contents, a natural move is to supplement the resources of an externalist theory with relations between thoughts or symbols. The result is a *long-armed conceptual-role theory*<sup>11</sup>—a theory that holds that a symbol's content is determined by its role in a thinker's mental economy, where that role is understood in a way that includes relations with the external world. As thus defined, the category includes many recent theories.<sup>12</sup> Such theories are external-internal hybrids.

It is also possible to maintain a purely internalist position, according to which a symbol's content is determined entirely by its relations to other symbols.<sup>13</sup> Such theories are *short-armed conceptual-role theories*. Thus, the term “conceptual-role theory” is used to encompass both purely internalist theories and theories that are external-internal hybrids. (As I explain in section III.B. below, however, restricting the term to a special subclass of such theories better captures what conceptual-role theorists intend.)

### III. The Proposed Taxonomy

#### A. *The Constitutive Question*

In this section (III.A.), I introduce and compare the organizing principles of the traditional and proposed taxonomies. In sections III.B. and III.C., I lay out the new taxonomy.

The traditional taxonomy divides up theories of content according to their implications for a non-redundant supervenience base for content facts.<sup>14</sup> (An easy way to see that the traditional classification works in this way is to notice how utterly standard it is to characterize the externalism/internalism divide in terms of supervenience.<sup>15</sup>) This characterization of the traditional taxonomy is underspecified, as the same phenomenon can have different supervenience bases at different levels, for example physiological and microphysical. Roughly, traditional taxonomists classify a theory of content by (its implications for) the most ontologically basic supervenience base that is mentioned by the theory being classified. I don't mean that traditional taxonomists always do so *under that description*. In fact, there are

two other roughly extensionally equivalent descriptions of the basis for classification that may in some cases better capture what theorists have in mind: 1) the most ontologically basic modal determinants of content that are *also* constitutive determinants of content; 2) the most ontologically basic modal determinants of content that can confidently and non-tendentiously be specified. (Shortly, I'll offer a diagnosis of how the traditional taxonomy comes to classify in this way.)

Let's use the term *basic modal determinants* for the elements of the lowest-level supervenience base mentioned by a theory of content. In these terms, the traditional taxonomy takes the organizing question to be *the basic modal determinants question*, or for short, *the modal question*: what are (a theory's implications for) the basic modal determinants of a thought's having a particular content?<sup>16</sup>

This formulation of the modal question must be understood as an attempt to characterize something that is ineliminably imprecise. For one thing, both first-order theorists and taxonomists do not always keep the modal/constitutive distinction clearly in focus.<sup>17</sup> And as suggested, some taxonomists probably frame the relevant question in one of two other roughly extensionally equivalent ways. For the sake of argument, we can assume that basic modal determinants are also constitutive determinants, for this assumption gives the best possible case to proponents of the traditional taxonomy (because it minimizes its differences from the proposed taxonomy). As we will see, the important differences between the traditional and proposed taxonomies are driven by the traditional taxonomy's classification by non-redundant supervenience base; whether or not the elements of the supervenience base are also constitutive determinants will not be the issue.

My proposed taxonomy rejects the modal question as the fundamental organizing question, in favor of *the constitutive question*: what (according to the theory under consideration) makes it the case that a symbol has a particular content (or expresses a particular concept)?

In order to get clear about why the difference between the modal and constitutive questions matters, we need to say more about constitutive accounts.<sup>18</sup> We can distinguish different kinds of constitutive accounts, with different ambitions.

A *full* constitutive account of phenomenon X purports to give an account of what X *is*. In the case of mental content, a full constitutive account purports to say for an arbitrary content what it is for a mental state to have that content.

More modestly, a theorist who denies that it is possible to give a general account of what X *is* may still try to say, for particular cases, what makes the X facts obtain. (It will be convenient to talk in terms of facts, but the discussion could be rephrased in terms of properties.) For example, even if we cannot say in other terms what it *is* to have an obligation, we might be

able, in every case of an obligation, to give an explanatory account of what it is in virtue of which the obligation obtains. One person may have a certain obligation because she has sworn an oath, another because he has knowingly participated in a practice and accepted its benefits, and so on. Similarly, a theorist could specify different mechanisms in virtue of which thoughts have their contents, without being able to say what unifies these various content-making mechanisms.<sup>19</sup> One thinker may have a concept of a particular substance in virtue of a causal chain that links her to instances of the substance, and another thinker in virtue of deference to an expert.

There are also projects intermediate between such a *specification of X-making features* and a full account of what X is. An account may purport to say something about what unifies the various X-making features but without offering a full account of what X is.

Another dimension of variation concerns the extent to which explanatory circles are permitted. A more ambitious kind of account requires that the facts to which the account appeals be explanatorily more basic than the target facts. A less ambitious account allows appeal to facts that are explanatorily co-equal with the target facts. We can use the term *reductive* for an account that is both full and non-circular. I use the term “constitutive account” to encompass the range of different projects just sketched, rather than reserving it for, for example, reductive accounts.

Many philosophers would recognize that a constitutive account seeks to provide something that cannot be captured in purely modal terms.<sup>20</sup> For they are aware of examples that show that there are modally necessary and sufficient conditions for being an X that do not figure in an account of what makes something an X. I offered the triangle example in the Introduction. To take a different kind of case, suppose that we could specify in microphysical terms necessary and sufficient conditions for an event’s being a murder. Such conditions plausibly do not figure in a constitutive account of murder.

It is very natural to take the moral of such examples to be that constitutive determination is a stronger relation than modal determination—i.e., that anything that is a constitutive determinant is a modal determinant, but not the converse. And given that view, it is natural to conclude that nothing will be lost if an account of a phenomenon is given in the form of a supervenience base, as long as everything in the supervenience base is a constitutive determinant.

I can now offer a speculative diagnosis of why the traditional taxonomy ends up classifying by basic modal determinants. In the first place, modal notions, such as supervenience, are more familiar and seemingly more precisely definable than notions such as the constitutive or the essential. As a result, when doing metaphysics, philosophers tend to reach for modal notions.

Next, once it is assumed that the goal is to specify modal determinants—a supervenience base—a theorist will want to specify a

*non-redundant* supervenience base, i.e., one such that no elements can be removed while preserving supervenience. To the extent that the project is understood in modal terms, modally redundant elements are doing no work. Now for a given phenomenon, there may be more than one non-redundant supervenience base, at different levels. Since the project is a metaphysical one, it is natural to focus on the supervenience base that is at the most ontologically basic level available.

It is easy to see how this could lead to a focus on the modal determinants of content at the lowest level specified by the theory being classified. For it will often be difficult to be confident about modal determinants below that level. And it would be tendentious to classify a theory on the basis of views (about lower-level modal determinants) that the proponent of the theory may not share. Since theories of content seek to give constitutive accounts, the lowest-level determinants mentioned by a theory of content will in general be constitutive determinants (or ones taken to be constitutive by the theorist). Moreover, as just sketched, there is a superficially plausible line of thought according to which nothing will be lost by working in modal terms—by formulating an account as a supervenience base or as modally necessary and sufficient conditions—as long as the elements to which the account appeals are constitutive as well as modal determinants. This line of thought may satisfy taxonomists who have an inchoate recognition that what is relevant are constitutive rather than merely modal determinants.

The foregoing is not meant to be a demonstration that the modal question is the organizing principle of the traditional taxonomy, but merely a diagnosis of how it comes to be that principle. The main evidence that the traditional taxonomy is based on the modal question will be the ways in which the proposed taxonomy differs from the traditional taxonomy.

We can now turn to the question of how classification based on the constitutive question differs from classification based on the modal question. I will show that there are elements of a constitutive account of a phenomenon that will have no place in a specification of basic modal determinants.

First, it is intuitively plausible that particular necessary truths may figure essentially in constitutive accounts of some phenomena. For example, on one well-known kind of view, legal norms obtain partly in virtue of fundamental moral or other normative truths, though it is plausible that such truths are necessary truths.<sup>21</sup> Whether or not this view is correct, it certainly seems to be coherent. But a specification of a (non-redundant) supervenience base can give no role to necessary truths. For necessary truths do not figure non-vacuously in the supervenience base for any domain. This point is particularly important for purposes of this paper since I will focus on the role of normative truths in constitutive accounts of content.

Second, a related point is that domains of necessary truths have no non-vacuous supervenience bases. Thus, if there can be any constitutive account

of fundamental normative truths, the elements of such an account cannot be captured by the specification of a supervenience base.

Third, a theory of a phenomenon may appeal to constitutive determinants at an intermediate level, and then go on to give an account, whether constitutive or merely modal, of those intermediate facts at another level. From a modal point of view, at least one of the two levels of determinants is typically redundant. Therefore, if we classify on the basis of the modal question, we will exclude intermediate layers of constitutive determinants from consideration.

Fourth, in particular cases, brute metaphysical necessities may prevent constitutive determinants from making an appearance in a supervenience base. For example, the constitutive dependence of the mental on features of the thinker's environment is consistent with the supervenience of the mental on intrinsic properties of the thinker.<sup>22</sup> Suppose that differences in the external world metaphysically necessitated appropriate differences in non-relational, physical properties of the thinker's brain. In that case, local supervenience could hold, even though the mental was constitutively dependent on the external world.<sup>23</sup>

In the next two sections, we'll see how the differences between the constitutive and modal questions lead to perhaps surprising differences in classification. One terminological point. In common philosophical usage, phrases such as "in virtue of which," "makes it the case that," and "determines" are used variously, sometimes to express purely modal notions, sometimes to express constitutive or other notions. I will use "in virtue of which" and "makes it the case that" as my canonical way of expressing constitutive accounts.

## B. *Non-Normative Theories*

*Non-Normative Versus Normative Theories.* One very general type of answer to the constitutive question is that what makes it the case that a symbol has a certain content is that certain non-normative facts obtain. Theories that give this type of answer are *non-normative theories* of content. In contrast, *normative theories* maintain that what makes it the case that a symbol has a particular content is at least in part that certain normative facts obtain.<sup>24</sup> Assuming that normative truths supervene on relatively high-level non-normative facts, the distinction between normative and non-normative theories will not in general be visible in the traditional taxonomy, to the extent that this taxonomy classifies by basic modal determinants.

I discuss non-normative theories in this section and normative theories in the following one. On the non-normative side, an important class of theories will be those that hold that what it *is* for a symbol to have a certain content is for the thinker to have a certain disposition to use that symbol. (As I use the term "disposition," what dispositions one has is a matter of



empirical and non-normative fact.) A *dispositional theory*, so understood, is not merely any theory on which dispositions are constitutive (or modal) determinants of content. (As we will see, a theory on which dispositions are constitutive determinants of content may even be a normative theory.) Rather, such a theory gives a reductive account of having content in terms of dispositions.<sup>25</sup>

Armed with this understanding of a dispositional theory, we can turn to the new taxonomy's treatment of indication and conceptual-role theories. I will argue that the new taxonomy does a better job at classifying the various theories that the traditional taxonomy puts in these categories and at capturing what is distinctive about the theories. In order to make this argument, I need to describe a special subclass of dispositional theories that I call *Mastery Theories*.

As a preliminary matter, let us distinguish the question of what concepts or contents are, or how they are individuated, from the question that theories of content seek to answer—the constitutive question of what makes it the case that a symbol or thought has a particular content. An answer to the constitutive question is consistent with a variety of different answers to the former *concept-individuation* question. Mastery Theories have in common a distinctive, though abstract and general, answer to the constitutive question.

***The Mastery View:*** what makes it the case that a symbol expresses a given concept is that the thinker has a disposition (whether successfully exercised or not) to use the symbol in accordance with the concept's *canonical pattern of use* (as explained below). The Mastery View should be understood as an attempt to say what having a content is, not merely to give a specification of content-making features. In intuitive terms, the Mastery View maintains that for a thought's content to involve a given concept is for the thinker to exercise his or her mastery of that concept in having the thought.

On any of a wide range of views about the individuation of concepts (as opposed to views about what it is to have a concept), each concept will be associated with a distinctive pattern of use of a symbol, either because that pattern of use itself individuates the concept or because there is a straight-forward mapping from what individuates the concept to the pattern of use. The pattern of use that is associated with a concept by the theory of concept individuation can be called the concept's *canonical* pattern of use. The natural, though not especially precise, idea behind the Mastery View is then that what makes it the case that a thought involves a particular concept is the thinker's exercising mastery of the concept, which is a disposition to use the concept in conformity with its canonical pattern.<sup>26</sup>

Evidently, positions on canonical patterns of use, and therefore on concept mastery, will reflect positions on concept individuation. For example, if one thinks that concepts are individuated by their relations to the

world—by the objects or properties to which they refer—one will take each concept's canonical pattern of use to be, roughly, a pattern of covariation with, or application to, instances of the relevant property. If, on the contrary, one thinks that concepts are individuated by their relations to other concepts, one will take each concept's canonical pattern of use to be the corresponding transitions between symbols or thoughts.

With this explanation of Mastery Theories and the Mastery View, we are ready to return to indication and conceptual-role theories. Let's begin with indication theories. The traditional taxonomy classifies them as theories that hold that a symbol's content depends on the thinker's dispositions to use the symbol with respect to the external world. The new taxonomy classifies indication theories among Mastery Theories.

In order to see that indication theories accept the Mastery View, we need to begin with their assumptions about concept individuation. Indication theories of content assume, roughly speaking, that concepts are individuated by the objects or properties to which they refer. Given this view about concepts, the Mastery View would say that what makes it the case that a symbol expresses a certain concept is that the thinker has a disposition to apply the symbol to instances of the concept. And this is exactly the position of an indication theory.<sup>27</sup>

Thus, an indication theory is a special kind of dispositional theory, a Mastery Theory. The traditional taxonomy's classification of an indication theory as an externalist theory on which dispositions, as opposed to history, determine content does not do justice to what is distinctive about such a theory.

The situation is more complex with respect to theories that, if classified in accordance with their implications for the basic determinants of content, would fall in the conceptual-role category. Most of these theories are Mastery Theories, as I will discuss shortly. One exception is Brandom's (1994; 2000) view, according to which what makes it the case that a thought has a certain content is that the thinker has certain inferential commitments and entitlements—where the existence of a commitment or entitlement is a normative fact. The (non-normative) dispositions of the thinker and of others in her community are relevant, Brandom thinks, because they are part of an explanation of the relevant normative facts. But since the claim is not that having a concept is having a disposition, the theory is not even a dispositional theory (in the above-defined sense), let alone a Mastery Theory. I'll say more about Brandom's view when I discuss normative theories in section III.C.

Similarly, if we classified theories solely by their implications for the basic determinants of content, Tyler Burge's position would be classified as a (long-armed) conceptual-role theory. Burge (1979; 1982; 1986) has argued that the constitutive determinants of content include facts about the

thinker's social and physical environment as well as facts about the inferences or judgments the thinker makes or is disposed to make.

At least *prima facie*, a dispositional theory, and especially a Mastery Theory, is not well suited to Burge's position. First, Burge does not attempt to offer a full account of what it *is* to have a content, but instead gives a specification of the features in virtue of which various kinds of thoughts have their contents.<sup>28</sup> Second, he holds that thinkers who have incomplete understanding of a concept, whether partial or incorrect, can nonetheless have thoughts involving it. The relevant symbol-using dispositions of a thinker with (a particular) incomplete understanding of a concept will be very different both from those of thinkers who fully grasp the concept and from those of thinkers with *different* incomplete understandings of the concept. Perhaps some nuanced non-Mastery dispositional theory could be consistent with Burge's views on incomplete understanding. But a normative theory of content would seem to fit those views better because such a theory can easily and naturally accommodate cases of incomplete understanding. In brief, the reason is that the relevant normative facts may be the same with respect to thinkers who have very different dispositions. For example, thinkers can be subject to the same standard, though some are disposed to satisfy the standard and others are not.

Paul Horwich's (1998) "use theory of meaning"<sup>29</sup> qualifies on the traditional taxonomy as a conceptual-role theory since it maintains that uses of symbols, including internal ones, determine content. But Horwich maintains that the relation between the determinants of meaning and meaning need not be the intuitive one that figures in the Mastery View (or any other intuitive relation) (1998, 65–71). Indeed, Horwich maintains that it is not possible to give a general account, for an arbitrary content, of what it is for a symbol to have that content.<sup>30</sup> Hence according to the new taxonomy, Horwich's position should not be lumped with conceptual-role theories.

As noted, most theories that the traditional taxonomy classifies as conceptual-role theories are Mastery Theories. Again, we need to begin with concept individuation. Conceptual-role theories assume, roughly speaking, that concepts are individuated by connections between propositions. (For simplicity, I'll focus on short-armed conceptual-role theories here, but corresponding points apply to long-armed conceptual-role theories as well.) Given this view of concept individuation, a Mastery Theory would say that what makes it the case that a symbol expresses a certain concept is that the thinker has a disposition (with respect to that symbol) to make the transitions that correspond to the concept's individuating connections. And this is exactly the position of most theories traditionally placed in the conceptual-role category. I will hereafter reserve the term "conceptual-role theory" for those theories within the traditional conceptual-role classification that are Mastery Theories (since in my view, these theories are what theorists most frequently have in mind when they discuss conceptual-role theories).

I have argued that indication and conceptual-role theories share a fundamental answer to the constitutive question. On both kinds of theories, a thought's having a certain content is explained in terms of the thinker's exercising mastery of the relevant concepts.

This commonality between indication and conceptual-role theories is invisible to the traditional taxonomy, as a result of its exclusive focus on basic determinants of content. The basic determinants of content are a function of at least two elements: a position on how concepts are individuated and a view about what it is to have a concept. Because of their different assumptions about how concepts are individuated, indication and conceptual-role theories have very different implications for the basic determinants of content. At the level of basic determinants, there is no way to separate out the contributions of different elements of the theories. The traditional taxonomy therefore misses the deep connection between indication and conceptual-role theories.<sup>31</sup>

The claim is not that a difference in basic determinants is irrelevant or should be ignored. The proposed taxonomy can reflect the differences between indication and conceptual-role theories, for example by subclassifying Mastery Theories by their positions on concept individuation or by their implications for basic determinants. By contrast, a classification based on basic determinants of content has no way to recognize aspects of the space of possibilities that are not straightforwardly reflected in basic determinants of content.

It is worth noting that the general point about the traditional taxonomy's shortcomings is independent of whether I am right about the particular case of indication and conceptual-role theories. The point holds as long as there are cases in which the structure of theories of content will not be recoverable from a specification of basic determinants. If we formulate a theory in terms of its implications for the basic determinants of content, we collapse the fine-grained structure of the theory into a single layer. In the present case, this has the effect of obscuring the existence of an important common element in two prominent kinds of theories. In other cases, we will see that it has the effect of obscuring important differences between theories.

Finally in this section, what about causal-chain theories? A pure causal-chain theory that purported to give a full account of having content would not accept the Mastery View. A thinker's acquisition of a symbol could trace back through a causal-chain to an object or property, though the thinker is not disposed to apply the symbol to that object or property. So the fact that a given symbol is connected by a causal chain to a particular object or property is not an example of one's being disposed to use the symbol in the canonical pattern of a concept with that reference. Indeed, such a pure causal-chain account need not be any kind of dispositional theory.

How a theory with causal-chain elements should be classified obviously depends on the rest of the account. A causal-chain story might even be part of a normative theory of content. For example, it might be that what makes a symbol have a certain content is that the thinker is subject to a requirement not to apply the symbol to anything other than the object or substance that lies at the end of a certain causal chain. (A theory that made such a claim could go on to explain what it is in virtue of which thinkers are subject to such requirements.)

### *C. Normative Theories*

As noted, the new taxonomy makes visible a family of theories of content that is invisible as such on the traditional taxonomy. *Normative* theories of content, in the sense in which I use the term, hold that what makes it the case that a symbol has a certain content is, at least in part,<sup>32</sup> the obtaining of certain normative facts. Normative facts include facts about what is correct and incorrect, better and worse, proper and improper, ideal and defective; and about what standards or requirements people are subject to.<sup>33</sup>

In order to make the discussion concrete, consider a particular group of normative theories—normative counterparts to Mastery Theories. According to a Mastery Theory (see section III.B. above), what it is for a symbol to express a concept *C* is for the thinker to be disposed to use the symbol in *C*'s canonical pattern of use—to be disposed to  $\Phi$ , let us say. Now compare a theory that holds that what it is for a symbol to express *C* is for the thinker to be subject to a standard requiring her to  $\Phi$ . This normative theory differs from the corresponding Mastery Theory only in substituting what the thinker is *supposed* to do for what the thinker is *disposed* to do. It is the fact that the thinker is supposed to use the symbol in a certain way, rather than the fact that she is disposed to do so, that makes it the case that the symbol has the appropriate content.

There are a variety of ways of developing such a normative counterpart to a Mastery Theory. For example, a normative theory need not maintain that normative facts specific to each token thought determine that thought's content – an approach that encounters obvious problems. An alternative is to maintain that it is abilities or dispositions, not individual thoughts, that are normatively individuated in the first instance. In other words, a particular disposition is, say, the thinker's possession of the concept of addition rather than of the concept of quaddition (Kripke 1982), not because of the empirical facts about the disposition but because of what the disposition is *for* or how it is *supposed* to be. And it is because the thinker is exercising that disposition that a particular thought's content involves the concept of addition.

What is crucial to normative theories of content is that normative facts are among the constitutive determinants of content facts. The relevant normative facts are not merely the consequence of content facts.

It is no part of a normative theory of content to claim that the relevant normative facts cannot be further accounted for in non-normative terms. (Contrast, for example, the Kripkenstein idea that the relation of meaning to use is normative in a way that rules out any attempt to give a constitutive account of content in terms of dispositions (Kripke 1982, 24, 37).) As we will see, a normative theory of content is consistent with the possibility that a constitutive account of the normative facts themselves can be given in non-normative terms. A distinction drawn in section III.A. between different kinds of constitutive accounts is pertinent here. Suppose the relevant type of normative fact is that a thinker is subject to a requirement or standard. One possibility is that we can give a full account in wholly non-normative terms of what it is to be subject to a standard. (Normative theories that give a reduction of the operative normative facts to non-normative ones raise interesting issues for the proposed taxonomy, but space does not permit discussion.) Short of such a reductive account, however, it may be possible to give a specification of non-normative features in virtue of which thinkers are subject to standards.

We have encountered some examples of normative theories. Brandom (1984) holds that symbols have their contents in virtue of normative facts—facts about inferential commitments and entitlements. In his terminology, the “proprieties” of linguistic practice, which are explicated in terms of commitments and entitlements, “confer” contents (159). He goes on to give a complex account of the relevant normative facts in terms of practices, including attitudes of taking people to be committed or entitled (166). Brandom is clear, however, that his account of the normative facts does not purport to reduce them entirely to non-normative facts. Rather, his account is normative “all the way down” (39, 45–46, 638, 648–649).

Certain theories of content that appeal to biological facts are normative theories. According to one kind of theory, what makes it the case that a symbol has a certain content is the obtaining of certain facts about function. (The relevant function can be that of the symbol, the mechanism that produces the symbol, or the mechanism that uses it.) The theories in question understand the notion of function normatively: for something to have a function is not for it to behave in a certain way, but for it to be correct or proper for it to do so, or defective for it to fail to do so.<sup>34</sup> For example, on this understanding, to say that an immune system’s function is to prevent disease is to say that a correctly operating immune system prevents disease—that that is what it is *for*.

We can use the term *teleological* theories of content for normative theories according to which the relevant normative facts are facts about function. Such theories can go on to give a constitutive account of function.

A *biological* teleological theory, then, is a teleological theory of content that maintains that biological facts—typically facts about evolutionary history—make it the case that the relevant function facts obtain.<sup>35</sup>

A biological teleological theory can usefully be contrasted with another kind of theory of content that appeals to biological facts. According to such a theory, a mental state's having a certain content is simply constituted by certain facts about the thinker's biology or physiology. Normative facts do not enter the picture. Paul Churchland's (2005) and Searle's (1992) positions are possible examples.

The new taxonomy takes Brandom's theory and teleological theories to be members of a family of normative theories. Members of the family vary with respect to which normative facts play a role in determining content, the role that those normative facts play, the source of those facts, and the extent to which they can be explained in terms of non-normative facts. Recognizing the family relations between normative theories helps to make clear that a normative theory's position on what it is for a thought to have a particular content is separable from its position on the sources of the relevant normative facts. For example, one could accept a biological teleological theory's account of content in terms of function, while substituting a different account of function, perhaps a non-biological one.<sup>36</sup> This point is an example of one that is obscured by understanding a theory as a specification of basic determinants.

A normative theory of content has the potential for a kind of explanation of the basic determinants of content. This claim requires some elaboration. It is natural to explain modal facts in terms of constitutive ones.<sup>37</sup> An instance is the familiar point that the supervenience of one domain on another is susceptible of different ontological (or other) explanations. For example, mereological relations can explain supervenience (Kim 1993, 156–159; 165–168).

A related idea is that we can appeal to a constitutive account of a phenomenon to explain why it has a particular supervenience base. Whether an event is a regicide does not supervene on intrinsic properties of the event. Here is an explanation of this fact: what makes an event a regicide is that it is a killing of a king, and whether a person is a king depends on complex facts about the person's relations to others. In some cases, we can appeal to a normative constitutive account to explain why a phenomenon has a particular supervenience base. Whether an event is a murder does not supervene on intrinsic properties of the event. Part of the explanation of this fact is that what makes a killing a murder is, among other things, its being unjustified, and whether a killing is unjustified depends on complex relational facts, such as whether the killing is carried out by a soldier participating in a war or by an executioner carrying out a sentence.

In the case of theories of content, an appeal to normative facts at an intermediate level can provide a kind of explanation of the basic

determinants of content. We have some prior understanding of the kinds of non-normative features in virtue of which normative facts obtain. A normative theory can appeal to that understanding to explain basic determinants of content. Thus, Brandom's appeal to dispositions and attitudes will be plausible to the extent that those are the right sort of raw materials to yield commitments and entitlements.

On biological teleological theories, in contrast to non-normative historical or biological theories, certain historical facts are alleged to be determinants of content *because* they are a source of function. So biological teleological theories have a distinctive explanation of why historical facts are basic determinants of content.

Similarly, many theorists accept that a person who incompletely grasps a concept can have thoughts involving the concept in virtue of deference to other people who fully grasp it. Since deference is plausibly a mechanism that can make an agent subject to a requirement, a normative theory of content has the potential to explain what having content is such that deferring to others is a way of accomplishing it.<sup>38</sup> More generally, a normative theory of content can explain what various content-making features have in common in virtue of which they are all mechanisms for having a particular content. It is worth noting that this kind of explanation may be available even if a reductive account of content is not possible.

#### IV. Observations and Comparisons

In this concluding section, I bring together a number of brief observations about the two taxonomies.

1. Since theories with very different structures can have the same implications for basic determinants of content, the traditional taxonomy tends to collapse different kinds of views into one basket. For example, a wide variety of views are classified as dispositional theories. The taxonomy obscures the fact that dispositions may play very different roles in different theories.

2. Conversely, the traditional taxonomy misses important commonalities between theories. Conceptual-role and indication theories share a general view about what having content is. But this central common feature is lost when the theories are formulated as specifications of basic determinants of content. As a result, conceptual-role and indication theories are often seen as polar opposites in the debate.

3. The new map gives an appropriate location to certain theories that sit uncomfortably in the traditional map. Biological teleological theories are a good example. The traditional taxonomy classifies such theories as a variation on externalist themes. But such theories' positions on what it is



to have content are fundamentally opposed to the positions of, for example, dispositional and (non-normative) biological theories.

4. More generally, the new taxonomy makes visible a family of theories of content that share a distinctive approach. Normative theories of content maintain that normative facts are (part of) what makes it the case that a symbol has its content. Because normative facts are modally redundant, the traditional taxonomy has no way of treating such theories as distinctive.

5. Once we stop conceiving of particular normative theories in terms of their implications for the basic determinants of content, we see that their positions on what it is to have content are separable from their positions on the source of the relevant normative facts. For example, it is possible to accept a teleological theory's account of content in terms of function, while rejecting its account of function.

6. A normative theory of content can take a variety of positions on the issue of whether the relevant normative facts can be accounted for in non-normative terms. For example, a theory can offer a reduction of those normative facts to non-normative ones. Alternatively, a theory can try to specify non-normative facts in virtue of which the normative facts obtain, but without attempting to say in non-normative terms what it is for the relevant normative facts to obtain.

7. A normative theory of content has the potential to offer a kind of explanation of constitutive determinants of content. For example, on Brandom's view, dispositions are determinants of content because they are a source of entitlements and commitments. Similarly, a normative theory may be able to explain what diverse content-making features have in common. For example, it may be that deference to other people is one of several ways to have a concept because it is one of several mechanisms that can make a thinker subject to a requirement.

8. Finally, the comparison of the two taxonomies illustrates the importance of distinguishing a constitutive account of a phenomenon from a specification of constitutive determinants that make up a supervenience base for the phenomenon. The point goes beyond the distinction between constitutive and modal determinants. Essential features of a constitutive account are not in general preserved when the account is formulated as a specification of basic determinants.

## Notes

1. I am indebted to Joseph Almog, Paul Boghossian, Tyler Burge, Martin Davies, Gilbert Harman, Barbara Herman, Pamela Hieronymi, Paul Horwich, Susan Hurley, Ram Neta, Philip Pettit, Jim Pryor, Gideon Rosen, Seana Shiffrin, Ori Simchen, and Jeff Speaks for helpful discussions or detailed comments on a draft. I would like to thank the participants in the January 2005 SOFIA conference, especially the organizers Ernie Sosa and Enrique Villanueva. I am very

grateful to Scott Shapiro and David Sosa, my commentators at that conference, for their thought-provoking papers. I would also like to thank the participants in the University of California, Irvine philosophy colloquium. Special thanks to Philip Pettit for encouraging me to turn these ideas into a paper, and to Martin Davies for many invaluable conversations.

2. It is confusing to talk of a “theory of content” since the theories in question are not theories of what contents are, but theories of having content—of what makes it the case that thoughts have particular contents. (See section III.B. below). Nonetheless, I’ll follow the familiar practice of using “theory of content” to mean a theory of having content.
3. Thanks to Julie Roskies Litman for this example.
4. Much of what I say will apply to theories of linguistic content as well.
5. See e.g., Rey 1997; Fodor 1987; 1994. There are, of course, other ways of classifying the aspects of a symbol’s use or causal role that are supposed to determine its content. For example, theories can rely on communal rather than individual use, or on optimal rather than actual use (see Boghossian 1989). Another variable is how sparse the specification of use is required to be. At one end of the spectrum, the relevant use must be specified in purely causal or physical terms (e.g., Dretske 2000; Fodor 1987; 1990). At the other end, mental notions can be used (e.g., Harman 1987; Peacocke 1992).
6. Indication theories focus on symbols’ relations to the world on the “input” or perception side. Another possibility is to appeal to the “output” or action side (Stalnaker 1987). A pure externalist theory cannot avail itself of this option, however. The problem is that the connection between a representation of the world—a belief or judgment for example—and an action depends on another internal state, such as a desire. The appeal to action-side symbol-world relations thus fits more comfortably in a long-armed conceptual-role theory, discussed in the text below.
7. E.g., Stampe 1977; Fodor 1980/1990. Dretske (1981) and Stalnaker (1987) offer non-biological versions.
8. See Ruth Millikan (1984; 1993), Karen Neander (1991; 1995), and David Papineau (1987; 1993).
9. E.g., Rey 1997, 243–249; Boghossian, 1989, 537. For an example of a writer expressing such uncertainty, see Boghossian 1989, 537, fn. 50 (“I shy away from saying whether R. Millikan . . . presents a theory of this form.”).
10. See, e.g., Devitt 1981; Devitt and Sterelny 1999.
11. Some writers reserve the term *conceptual-role theory* for a position that holds that the content of a symbol *is* the role of the symbol. My terminology does not saddle conceptual-role theories with this implausible position. Conceptual-role theories, like indication theories, hold that the role of a symbol *determines* its content; the two kinds of theories differ in which aspects of that role are relevant. For general discussion, see Greenberg and Harman (forthcoming).
12. E.g., Harman (1973; 1982; 1987), Peacocke (1992), Pettit (1993), Horwich (1998), Brandom (1994; 2000).
13. E.g., Field 1977; Loar 1981.
14. To say that the B facts supervene on the A facts is to say that the A facts (or properties) modally determine the B facts (or properties)—that there can be no

difference in the B facts without a difference in the A facts. Modal determination can be metaphysical determination or something else, such as nomic determination, depending on the type of possibility in terms of which it is defined.

15. This tendency goes back to the origins of externalism. Putnam (1975, 219, 270), though not crystal clear, seems to define externalism modally. Burge (2003, 302) notes that his earlier work suggested such a characterization. For other examples (chosen arbitrarily), see Jackson and Pettit (1993); Braddon-Mitchell and Jackson (1996, pp. 211, 216) and Brown (2005, 14).
16. The notion of a basic modal determinant is obviously theory-relative. The modal question should be understood as relativized to the theory being classified.
17. For an example of a theorist giving an apparently modal characterization of his project, see Fodor (1990, 96). Boghossian (1989, 532–34) addresses the issue of what should be required beyond modally necessary and sufficient conditions, and suggests a requirement of “intensional equivalence.”
18. I prefer “constitutive account” to “analysis” for at least two reasons. First, the latter is more likely to suggest an account of the meaning of a term. Second, “analysis” suggests a reductive account, while, as discussed in the text, a theorist can specify constitutive determinants of a phenomenon without offering a reduction. For a helpful discussion of constitutive accounts, see Rey (1987, 28–34). See also Millikan (1993, 16–17).
19. Thanks to Tyler Burge and Gideon Rosen for pushing me to clarify the distinction between full constitutive accounts and specifications of X-making features.
20. Kit Fine (1994) makes a powerful argument that the notion of essence that plays a central role in the metaphysics of identity cannot be captured in modal terms. See also Almog (1996). Readers may wonder how my constitutive/modal distinction relates to Fine’s essence/modality distinction. Although Fine’s argument is congenial to my position, I should register two differences between the kinds of accounts with which we are concerned. First, Fine’s discussion concerns accounts of the identity of particular objects (what Socrates is), while I am concerned with types (e.g., my question is *not* what it is in virtue of which something is a particular token thought, but what makes a token thought have the content (type) that it does). Second, Fine is concerned with reductive accounts (as defined in the text), while, as noted, I want to understand constitutive accounts broadly enough to include less ambitious kinds of accounts.
21. See generally Greenberg 2004. The fundamental normative facts are the most general ones that, along with ordinary empirical facts, explain the obtaining of specific normative facts.
22. Burge (2003, 371–372) makes a closely related point about the supervenience of the mental on the physical. Unlike Burge’s example, mine does not depend on the idea that the relevant physical properties could themselves be anti-individualistically individuated. My example therefore extends to the consistency of externalism with the supervenience of the mental on individually individuated properties of the thinker.
23. Martin Davies (1998, 327–329) discusses the relations between constitutive and modal externalism. For discussion of a case related to the one described in the text, see Hurley (1998, chapter 8).

24. There are other important contrasts that could be drawn between answers to the constitutive question. For example, there is the contrast between theories that maintain, and theories that deny, that translation or interpretation plays an essential role in what makes it the case that a thought has a certain content.
25. For the relevant sense of “reductive,” see section III.A. above.
26. In explicating the notion of a canonical pattern of use, I use the vague notion of a straightforward mapping. The discussion in the text should help to clarify the notion. For further discussion and explication, see Greenberg MS. The vagueness in the notion, I suggest, aptly parallels the play in what counts as a conceptual-role or indication theory of content.
27. Many proponents of Mastery Theories add a proviso about deference to other people, though their theories cannot adequately make sense of the role that such deference is supposed to play. See Greenberg MS and section III.C. below.
28. For the distinction invoked here, see text accompanying notes 18–19 above.
29. Horwich’s concern is with linguistic meaning rather than mental content.
30. Another possible position is that dispositions are modal, but not constitutive, determinants of content. Soames (1998) points out that dispositions or other non-intentional facts may modally determine content facts without *a priori* entailing them. He does not discuss constitutive accounts, but the general thrust of his discussion—he doesn’t consider the possibility of taking “determination” to be constitutive determination – suggests a lack of sympathy for the possibility of a constitutive account, indeed for the very idea of constitutive determination.
31. Perhaps because of the influence of the traditional taxonomy, the fact that Mastery Theories are only a special subclass of dispositional theories seems to be frequently overlooked. Indeed, the literature often proceeds as though the space of possibilities were limited to variations on the Mastery-View theme.
32. Normative theories can give a more or less significant role to non-normative facts as well.
33. Philosophers sometimes distinguish “prescriptive” norms, which are said to entail “ought” statements, from other norms. Prescriptive normativity is thought to be a stronger form of normativity than the thinner normativity sufficient for misrepresentation. I am dubious about the value of this distinction, but we can stipulate that normativity sufficient for misrepresentation is all that is needed for present purposes.
34. Several well-known theories of content appeal to a notion of function that is explicated biologically. See Ruth Millikan (1984; 1993), Karen Neander (1991; 1995), David Papineau (1987; 1993). It is sometimes disputed whether these theories in fact appeal to a normative notion of function. As I interpret them here, the theories work with a genuinely normative notion, but purport to account for the relevant normative facts in non-normative terms. (See also the previous footnote.) To avoid exegetical questions, readers who take actual teleological theorists not to be working with a normative notion of function can consider hypothetical theorists who maintain the views discussed in the text.
35. As noted above (section II), biological facts can also enter a theory of content as part of an account of optimal conditions. A theory that relies on biology in this way may have a normative component. An example is a theory that holds that a symbol has a particular content in virtue of what the thinker is *actually* disposed to do under certain *normatively* characterized conditions.

36. For a non-biological notion of function in moral philosophy, see Foot (2001).
37. For discussion, see Fine (1994) and Almog (1996).
38. I argue elsewhere (Greenberg MS) that leading theories of content lack an adequate account of the role of deference.

## References

- Almog, J. 1996: "The What and the How II: Reals and Might". *Noûs*, 30, 413–433.
- Boghossian, P. 1989: "The Rule-Following Considerations". *Mind*, 98, 507–549.
- Braddon-Mitchell, D. and Jackson, F. 1996: *Philosophy of Mind and Cognition*, Oxford: Blackwell.
- Brandom, R. 1994: *Making It Explicit*. Cambridge: Harvard University Press.
- 2000: *Articulating Reasons*. Cambridge: Harvard University Press.
- Brown, J. 2004: *Anti-Individualism and Knowledge*. Cambridge: MIT Press.
- Burge, T. 1979: "Individualism and the Mental". *Midwest Studies in Philosophy*, 4, 73–121.
- 1982: "Other Bodies," A. Woodfield, ed., *Thought and Object*. Oxford: Oxford University Press.
- 1986: "Intellectual Norms and Foundations of Mind," *Journal of Philosophy*, 83, 697–719.
- 2003: "Reply to Owens," in M. Hahn and B. Ramberg, eds., *Reflections and Replies*. Cambridge: MIT Press.
- Davies, M. 1998: "Externalism, Architecturalism, and Epistemic Warrant" in C. Wright, B. Smith, and C. MacDonald, eds., *Knowing Our Own Minds*. Oxford: Oxford University Press.
- Devitt, M. 1981: *Designation*. New York: Columbia University Press.
- Devitt M. and Sterelny, K. 1999: *Language and Reality* (2d ed.). Cambridge: MIT Press.
- Dretske, F. 1981: *Knowledge and the Flow of Information*. Cambridge: MIT Press.
- 2000: "If You Can't Make One, You Don't Know How It Works," in his *Perception, Knowledge, and Belief*. Cambridge: Cambridge University Press.
- Fine, 1994: "Essence and Modality," *Philosophical Perspectives*, 8, 1–16.
- Fodor, J. 1987: *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. MIT Press.
- 1980/1990 "Psychosemantics or Where Do Truth Conditions Come From?," in W. Lycan, ed., *Mind and Cognition*. Oxford: Blackwell (1<sup>st</sup> ed.), 312–338.
- 1990: *A Theory of Content and Other Essays*. Cambridge: MIT Press.
- Foot, P. 2001: *Natural Goodness*. Oxford: Oxford University Press.
- Greenberg, M. 2004: "How Facts Make Law," *Legal Theory*, 10, 157–198, to be reprinted in E. Villanueva, ed., *Social, Political, & Legal Philosophy*, 2 (2005) and in S. Hershovitz, ed., *Exploring Law's Empire* (forthcoming Oxford University Press).
- MS: "Incomplete Understanding, Deference, and the Content of Thought" (unpublished manuscript).
- Greenberg, M. and Harman, G. forthcoming: "Conceptual Role Semantics" in E. Lepore and B. Smith, eds., *Oxford Handbook of Philosophy of Language*, Oxford: Oxford University Press.
- Harman G. 1973: *Thought*. Princeton, NJ: Princeton University Press.
- 1982: "Conceptual Role Semantics," *Notre Dame Journal of Formal Logic*, 23, 242–56.
- 1987: "(Nonsolipsistic) Conceptual Role Semantics," in E. Lepore, ed., *New Directions in Semantics*. London: Academic Press.
- Horwich, P. 1998: *Meaning*. Oxford: Oxford University Press.
- Hurley, S. 1998: *Consciousness in Action*. Cambridge: Harvard University Press.

- Jackson, F. and Pettit, P. 1993: "Some Content Is Narrow," in J. Heil and A. Mele, eds., *Mental Causation*. Oxford: Oxford University Press.
- Kim, J. 1993: *Supervenience and Mind*, Cambridge: Cambridge University Press.
- Kripke, S. 1980: *Naming and Necessity*. Oxford: Blackwell.
- 1982: *Wittgenstein on Rules and Private Language*. Oxford: Blackwell.
- Millikan, R.: 1984: *Language, Thought, and Other Biological Categories: New Foundations for Realism*. Cambridge: MIT Press.
- 1993: *White Queen Psychology and Other Essays for Alice*. Cambridge: MIT Press.
- Neander, K. 1991: "The Teleological Notion of Function," *Australasian Journal of Philosophy*, 69, 454–468.
- 1995: "Misrepresenting and Malfunctioning," *Philosophical Studies*, 79, 109–141.
- Papineau, D. 1987: *Reality and Representation*. Oxford: Blackwell.
- 1993: *Philosophical Naturalism*. Oxford: Blackwell.
- Peacocke, C. 1992: *A Study of Concepts*. Cambridge: MIT Press.
- Pettit, P. 1993: *The Common Mind*. New York: Oxford University Press.
- Putnam, H. 1975: "The Meaning of 'Meaning,'" in his *Philosophical Papers, Volume 2: Mind, Language and Reality*. Cambridge: Cambridge University Press.
- Rey, G. 1997: *Contemporary Philosophy of Mind: A Contentiously Classical Approach*. Oxford: Blackwell.
- Soames, S. 1998: "Skepticism about Meaning: Indeterminacy, Normativity and the Rule-Following Paradox," in *Meaning and Reference, Canadian Journal of Philosophy*, Supp. Vol. 23.
- Stalnaker, R. 1987: *Inquiry*, Cambridge: MIT Press.
- Stampe, D. (1977): "Towards a Causal Theory of Linguistic Representation," in P. French, T. Uehling, and H. Wettstein, eds., *Midwest Studies in Philosophy*, University of Minnesota Press.