More on the Interactive Indexing Semantic Theory

John Dilworth
Department of Philosophy
Western Michigan University

Abstract   This article further explains and develops a recent, comprehensive semantic naturalization theory, namely the *interactive indexing* (II) theory as described in my 2008 *Minds and Machines* article "Semantic Naturalization via Interactive Perceptual Causality" (Vol. 18, pp. 527-546).   Folk views postulate a concrete intentional relation between cognitive states and the worldly states they are about. The II theory eliminates any such concrete intentionality, replacing it with purely causal relations based on the *interactive theory of perception* (ITP). But intentionality is preserved via purely abstract propositions about the world that *index*, or correlate with, appropriate cognitive states.

  Further reasons as to why intentionality must be abstract are provided, along with more details of an II-style account of representation, language use and propositional attitudes. All cognitive representation is explained in terms of classification or sorting dispositions indexed by appropriate propositions. The theory is also related to Fodor's *representational theory of mind* (RTM), with some surprisingly close parallels being found in spite of the purely dispositional basis of the II theory. In particular, Fodor's insistence that *thinking about* an item cannot be reduced to sorting dispositions is supported via a novel two-level account of cognition--upper level propositional attitudes involve significant intermediate processing of a broadly normative epistemic kind prior to the formation of sorting dispositions. To conclude, the weak intentional realism of the II theory--which makes intentional descriptions of the world *dispensable*--is related to Dennett's 'intentional stance' view, and distinguished from strong (indispensable) intentional realist views. II-style dispositions are also defended.

This article further explains and develops a recent, comprehensive semantic

naturalization theory, namely the *interactive indexing* (II) theory (Dilworth 2008a, 2009,

forthcoming). The theory may be introduced as follows.

To begin, Fodor famously observed that, since intentionality doesn't go deep enough to be a fundamental physical property, then if intentional aboutness is real, it must really be something else (Fodor 1987, p. 97). Consequently, on his view, a realist view of intentionality must also be a physically reductive one. However, it has been little noticed that there is another, non-reductive way of being an intentional realist. This involves denying that there are any *physically* real intentional relations--i.e., holding an eliminativist view of physical intentionality--while yet also preserving the reality of intentionality via assigning it instead to the purely abstract realm of the propositionally structured metatheory of logic and science. On such a view, intentionality lives on in ways such as in the purely abstract, intentional or representational relations of propositions to their worldly truth-conditions, in spite of its banishment from the concrete physical world.

This purely abstract view of intentionality means that, strictly speaking, abstract intentional structures themselves cannot be naturalized, any more than pure mathematical structures or normative logical principles could be naturalized. Nevertheless, this does not prevent broader naturalization attempts--generally concerning how concrete cognitive systems relate to abstract propositional structures--from being provided, and possibly even being successful. One such attempt, specifically directed at the naturalization of semantics--the *interactive indexing* or II theory--was recently defended, as noted above. The II theory links a purely causal, *interactive* theory of perception (ITP), with an account of how such causally based perceptual states could be correlated with, or *indexed* by, appropriate abstract propositions.

To be sure, as discussed in Dilworth (2008a), some, such as Fodor (1990) and Dretske (1981), might want more out of a semantic naturalization project. They might propose a *global semantic naturalization* (GSN) thesis, according to which *all* semantic concepts, including those of proposition, truth, logic and so on, should be reductively explained in physicalist terms. But a more modest, II-style naturalization project, which respects the broadly realistic logical framework of standard, naturalistic scientific theories, arguably is at least equally legitimate. This article further develops various aspects of the II theory. In particular, more defense of the abstractness of intentionality will be provided, along with further details of an II-style account of representation, language use and propositional attitudes.

1. Why Intentionality Is Abstract

The interactive indexing (II) theory claim that intentionality is purely abstract may initially seem unexpected. But there are some strong presumptive arguments in favor of it, such as the following. To begin, methodological simplicity dictates that, at least initially, we should not presume that intentionality could be both abstract *and* physically real, for in those circumstances there would be two fundamentally different kinds of intentionality to be explained. But if we can only choose one kind, the following indispensability argument suggests that it should be the abstract kind. The argument is that virtually everyone acknowledges that at least *truth-evaluable* content is of some

3

propositional kind, and therefore, on most views of propositions, it would be abstract. But, since many intentional cognitive states have truth-evaluable contents--such as in propositional-attitude cases of belief, which are widely viewed as relating cognitive states of belief to abstract propositions--it would seem that we are already committed to explaining such truth-evaluable states in terms of the abstract intentionality involved in the relevant propositions that are the contents of the relevant beliefs.

To be sure, there are those, such as Fodor, who regard propositional attitudes as attitudes toward mental representations whose content, even if it is truth-evaluable, is nevertheless to be explained in purely causal terms (Fodor 1990, 2008). In this manner Fodor is able to maintain consistency with his physically reductive view of intentionality. However, at this point a burden-of-proof argument may be invoked. Aside from such proponents of some kind of global semantic naturalism (GSN) project, as discussed in the introduction, the standard presumption is that propositions are indeed truth-conditionally based, abstract entities. Perhaps that standard presumption could be overturned, but so far Fodor, Dretske et al have not yet persuasively shown how it could be achieved, and so the burden of proof is on them to actually carry out their nascent program.

But in the meantime, we may reasonably presume that, since propositions are an integral ingredient in standard theories of propositional attitudes, abstract propositional kinds of intentionality are currently an *ineliminable* factor in our understanding of mental content. Consequently, methodological simplicity dictates that we should--at least initially or *pro tem*--presume that *all* intentionality is purely abstract. The II theory is perhaps one of the

first semantic naturalization theories to explicitly accept that presumption, and to fully incorporate it in a theory, via its account of the indexing of non-intentional, concrete cognitive states by intentionally based abstract propositions.

Here are some related reasons for taking intentionality to be abstract. To begin, note that the concept of intentionality or aboutness is closely related to the broadly logical or rational concepts of propositional content, truth and falsity, logical conseqence or entailment, and inconsistency. Arguably this is a tightly coherent group of broadly rational or logical concepts, so that if one is causally reducible, then so also must the others be. But the idea that purely logical concepts, such as those of truth or entailment, could be reduced to causal concepts is clearly an unacceptable one.

Also, arguably intentionality *cannot be pried apart* from this group of logical concepts, because, for instance, the concept of the truth of an empirical proposition cannot be understood without understanding what it is for a proposition to be *about* a worldly fact, and for it to *correctly describe* the fact that it is *about*. Such an abstract proposition also stands in integral logical relations to other propositions. So it seems that we have no choice but to reject a Fodor-style attempted reduction of intentional relations to causal relations. The concept of intentionality closely coheres with a group of broadly abstract, logical or rational concepts, and consequently it must be so treated in any satisfactory naturalization of semantics.

Here are some more details on the II theoretical approach to these issues. Informally put, the basic idea is to kick upstairs, into the purely abstract realm, all semantic and intentional concepts, and then to regard concrete cognitive states as only being *correlated with*, or *indexed by*--rather than *instantiating*--appropriate kinds of truth-evaluable propositional content. So the overall II theoretical strategy is of a divide-and-conquer kind. Folk views of intentionality postulate a concrete intentional relation between distinct physical states, namely cognitive states and the worldly states they are about. But according to the II theory, the physical states are actually related only in purely causal ways, namely, those conforming to the interactive theory of perception (ITP). That is one basic division of the II theory. The other division is embodied in the claim that intentionality is real, but confined to the purely abstract propositional realm.

On this broadly mathematical or *abstract indexing* approach to content, intentional descriptions of cognitive states function in much the same way as mathematical structures do when used in the natural sciences. Arguably both intentional and mathematical structures provide abstract models that are of assistance in explaining and systematizing concrete worldly objects and events. But they do so without those abstract models and properties being, strictly speaking, actually instantiated by, or being actual properties of, concrete objects or events.

On such an account, just as a four-pound weight does not actually instantiate a property of fourness, so also a cognitive state S that is nominally about a worldly object x does not actually instantiate a property of aboutness with respect to object x. Instead, as recently

argued in support of the II theory (Dilworth 2008a, 2009, forthcoming), an adequate causal theory of perception must provide appropriate concrete cognitive structures which are such that the relevant, causally defined cognitive state S *qualifies as being indexed by* an appropriate abstract proposition that is about object x. It is for this reason that the overall theory may be described as the *interactive indexing* (II) theory, in that it combines an *interactive* causal theory of perception (ITP) with a propositional *indexing* strategy, as just described.

2. How Indexing Works

On the II approach to semantic naturalization, a systematic account is to be provided of how causally salient factors in cognition enable the resulting concrete cognitive states to be adequately correlated with, or indexed by, appropriate propositions--in place of a Fodor-style attempt to reductively naturalize meaning itself in causal terms. For current purposes, cognition generally will be approached via singular perceptual cases. In the case of a singular proposition about a particular worldly object X, such as the proposition p that an object X is red, an adequate casual theory must provide causal factors sufficient to identify three salient indexing constraints involved in a perception of the color of object X.

The relevant three constraints are those most basic to identifying which singular proposition indexes the relevant perceptual state. In the case of the proposition p = 'X is

red', there are at least three basic constraints. The object X that proposition p is about

must be identifiable as one particular object that is distinct from other similar objects.

The property that proposition p attributes to X must also be identifiable. And third,

proposition p must have truth-conditions, sufficient to determine whether it is true or

false.

The II theory identifies three corresponding, but purely causal, constraints applying to the

relevant concrete perceptual state S itself, which together determine which proposition

indexes state S. First, a *particularity* constraint--the theory must be able to explain how

the perception can be of that particular object X, rather than just of some object or other

that is similar to X. Second, a *property* constraint--the theory must be able to explain how

some particular property F of X is perceived, such as its red color, rather than some other

property. And third, a *correctness* constraint must also be satisfied--the theory must be

able to explain how the relevant perception could be correct or incorrect. For example,

for an object X that is red, some explanation must be available as to how a person Z

could correctly perceive it as being red in some cases, and also what would be involved

in person Z instead misperceiving object X as being some other color, such as green.

These three constraints, of particularity, propertyhood and correctness, are minimum

conditions on any adequate account of truth-evaluable perception of the properties of

particular objects that are perceived.

The *interactive theory of perception* (ITP), as the perceptual component of the II theory,

satisfies these constraints in the following way. The basic claim of the interactive theory

of perception is that a cognitive system Z perceives an object X just in case X causes some sense-organ z_i of Z to cause Z to acquire some X-related disposition D. Then, more specifically, Z counts as perceiving some property F of X just in case Z is caused to acquire an *F-related classification disposition* toward X. For example, if Z, as a result of perceiving X, becomes disposed to move object X in with other red objects, but not with objects of some other color, that would provide some evidence that Z had acquired the relevant red-related classification disposition toward X. And finally, the ITP satisfies the correctness constraint, in that behavioral evidence as to whether Z correctly or incorrectly classifies object X would be available as part of the evidence as to how Z has classified X.

To address a potential criticism, there might be a concern that the concept of an X-related disposition, or of a disposition that is 'directed toward' an object X, surreptitiously re-introduces an intentional concept into the causal analysis. However, this is not so, because in evolutionary terms the interactive causality is patently purely causal or dispositional, and the concepts of X-relatedness, or of a disposition directed toward X, arguably are no more than informal descriptions of salient aspects of the interactive causality. The basic concept of X-relatedness for Z's disposition D adverts only to the causal effects that activation of disposition D would have upon X. For example, if prey Z, as a result of perceiving predator X, becomes disposed to flee from X, and via the activation of its disposition thereby deprives X of a meal, then Z has caused X to go hungry on this occasion--as a result of Z having been caused by X to perceive X--which X would not have been, had Z failed to thus perceive X. So the relevant causality in X-

related dispositions is as basic and non-intentional as anyone might desire. (See section 10 for further considerations).

3. Representation, Intentionality and Meaning

The concept of representation is closely allied with concepts of intentionality and meaningful intentional content. As a case in point, in perceptual cases, it is generally assumed that truth-evaluable perceptual content is a species of representational content, so that veridical cases of perception are cases in which a perceptual state correctly represents some way in which the world is. As a consequence of this close connection, a consistent interactive indexing (II) theory must regard all genuine cases of representation as being purely abstract, as with intentional content generally. So on the II view, an abstract proposition can directly represent the world, but strictly speaking, concrete perceptual states as such do not represent anything.

Initially this might seem to be an unexpected or counter-intuitive view. It might seem obvious that perceptual states, and also external representations such as pictures, can directly represent worldly items. However, recall that in linguistic cases, it is a commonplace that strictly speaking, indicative sentences do not themselves represent the world, but instead they *express propositions*, which in turn do directly represent the world. So the II theory is committed to (what could be described as) a *global propositional view* of content--all broadly representational kinds of meaningful content,

including broadly pictorial kinds (Dilworth 2008b, c) are to be understood on the same indirect, abstract propositional model that is already familiar from linguistic cases.

Indeed, this point is already implicit in the basic indexing model of the current II theory, as applied to paradigm perceptual cases on which the interactive theory of perception (ITP) itself is based. Another, perhaps more intuitive way to characterize the view is that according to it, *all meaning is expressive meaning*. Any concrete meaningful units--whether external representations such as linguistic sentences and pictures, or internal representations such as cognitive or perceptual states--have meaning, or are meaningful, in virtue of their expressing abstract, broadly propositional kinds of meaning.

Nevertheless, this is not to deny that, for instance, a picture that shows how a whole group of objects are spatially related to each other--such as is provided by a visual perceptual experience, a mental image, or an external picture--expresses an abstract structural model whose meaningful pictorial structure cannot be reduced to a series of *linguistically* expressible propositions describing how each element in the structure is related to some other element.

Thus the indexing view can accept that there are significant differences between linguistic and pictorial methods of meaningful representation and expression. Its claim is instead that insofar as a picture is meaningful or representational at all, it is so in virtue of expressing an abstract pictorial structure--and one that is truth-evaluable, insofar as it is intended to be a representation of some actual, pre-existing structure. So pictures

preserve the abstractness and truth-evaluability of propositional kinds of meaning, in spite of the differences in their expressive functioning from those of linguistic expressions. Consequently, it is legitimate for the II theory to hold that all literal or non-fictional representational meaning, whether of pictorial or linguistic kinds, is abstract and potentially truth-evaluable.

4. Interactive Indexing Extended to Cognition Generally

Perception is a kind of cognition, but not the only kind. I shall now outline how the interactive indexing (II) theory--which initially explains how specifically perceptual states indexically relate to propositions--could be extended to cover more generic kinds of cognition. The concrete non-mental representations discussed in the previous section, such as sentence tokens or pictures, will play a central role in the discussion.

The basic idea is that perception of such representations plays an analogous role in cognition generally to that of perception of the actual objects and properties that such representations are about. For example, if person P perceives that a book X is red, the basic II theory claims that person P's perceptual state S is indexable by the proposition 'X is red' in virtue of the book X causing P, via one or more of his sense-organs, to acquire a red-related disposition toward book X--such as a disposition to put book X on a shelf reserved for red books. Now suppose that person P is a competent hearer of English sentences, who hears person Q--who is pointing towards book X--say "that book is red".

Arguably person P could acquire similar, or even the same, red-related dispositions toward book X from seeing Q's pointing and hearing person Q's utterance of the token sentence "that book is red" as he previously had acquired from his own perception of X itself.

This is not a pure or paradigm case of a perceptually acquired disposition, because the item that caused it--primarily, person Q's utterance of the sentence "that book is red"--is not the primary item toward which the hearer P acquires a disposition. Nevertheless, it is a disposition that P acquires as a result of perception of Q's utterance, and it is a typical case of how direct perception of a token representation can be used to indirectly acquire a disposition toward the object X that is being talked about.

Another, simpler example involves a picture rather than a sentence. In browsing through an illustrated catalog of first editions of Hume's writings online, one might see a photograph of a copy X that causes one to become disposed to buy it. This is a disposition directed toward the book X itself--similar to a disposition one would acquire when directly viewing the book--yet the disposition was acquired indirectly through perception of the photograph rather than of the book itself.

Examples such as these also serve to provide more evidence of the fundamental simplification of central representational issues achievable by the II theory. Though there are many complex issues concerning linguistic and pictorial representation, in such everyday cases items count as successful representations of a relevant item X just in case

a perceiver of them acquires the same--or similar enough--dispositions toward the represented item X as would have been acquired through direct perception of X itself. In such cases, the concept of representation can be explained in terms of the role of the relevant, representationally acquired dispositions in ensuring the propositional indexicality of the resulting cognitive state of the perceiver. This is just one specialized kind of case in which representational concepts as applied to a cognitive system are to be explained in terms of their role in ensuring propositional indexing of such cognitive states.

This account of the functioning of non-mental representations can also readily be extended to cover communicative intentions. Since perceivers or consumers of representations can thus acquire dispositions toward objects not directly perceived, social practices involving the production and communication of such non-mental representations can be explained in terms of the potential dispositional benefits to such consumers. As a simple case, many perception experiments are set up so that the experimental subject communicates with the experimenters by some standardized representational means, such as pressing a buzzer when a red item is perceived. In such a case, person P's perceiving that the object on the screen is red consists in his becoming disposed to classify it as red. His communicative intentions primarily involve his thereby becoming disposed to cause the experimenters to become disposed to classify the object as red, and his instrumental representational means of carrying out this task consists in his becoming actively disposed to press the buzzer under those perceptual conditions.

Of course, this is a very quick summary of a large and complex area, but it should be enough to show the initial viability of the II approach to non-mental representations.

5. Representation, Indexing and Cognitive Development

The II theory attempts to explain all cognitive and perceptual processing of representations in a completely uniform manner. As an initial overview, the basic idea is that, for a singular perceptual state that relates a perceiver to a particular worldly object X, a concrete perceptual state is representational just in case it is propositionally indexable (henceforth just: indexable). As discussed in section 1, indexability is a matter of a concrete perceptual state being able to satisfy the three basic constraints on a representational perceptual state--a particularity constraint, a property constraint and a correctness constraint.

On this account, a perceiver Z counts as perceptually representing some property F of a particular X just in case Z is caused by X to acquire an *F-related classification disposition* toward X.   This is not to deny that there also exist cases of non-indexable and non-representational X-related perceptual dispositions, such as a disposition of a predator Z to eat prey X as a result of Z perceiving X. But such non-indexable dispositional states are to be distinguished from indexable or representational dispositional states, which the II theory regiments as cases specifically involving *classification* dispositions. Thus, in the broader evolutionary setting of the growth of

perceptually based cognition, indexable representational perceptual states are probably a later outgrowth of non-cognitive, purely action-oriented perceptual dispositions.

A significant advantage of this II approach to perceptual representation--and to cognitive representation in general--is that it can completely avoid the endless disputes over what cognitive representation is, and what inner mechanisms should qualify as being genuinely representational (Cummins 1996, Ramsey 2007). On the II approach, all that matters is that a resultant perceptual or cognitive state should involve classification dispositions, and hence be propositionally indexable, no matter what kinds of inner mechanisms of indication, modeling, computational symbolization, neural net activation vectors, and so on, actually produced the relevant dispositional state. So the II theory can cheerfully embrace an inclusive pluralism about cognitive representational mechanisms, as long as they produce appropriate classification dispositions toward a relevant object under appropriate conditions.

Thus the II theory can provide a thoroughly deflationary account of mental representation, which can deny that there is any key property in virtue of which a cognitive state represents a worldly state. Also, it must not be forgotten that strictly speaking, according to the II theory, mental states as such do not represent the world at all, since the theory eliminates rather than reduces concrete mental representation. On this account, the only genuine representations of the world are abstract propositional representations of it.

Another significant advantage of the II theory that it is potentially able to integrate well with a broadly developmental view of cognition. In particular, there is conceptual room in the II theory for a *two-stage* account of the development of representational or classificatory cognition. In the first stage or level, we can distinguish relatively low-level kinds of animal behavior that provide some minimal evidence of cognitive classification, such as a successful ability by some animal species to distinguish edible from poisonous plants. On the basis of this evidence we could reasonably impute rudimentary or low-level beliefs to their members, concerning the edibility or non-edibility of various plant species.

However, these lower-level abilities, in spite of their being minimally representational, and hence propositionally indexable in the relevant ways, do not automatically provide any evidence of more sophisticated cognitive abilities on the part of members of those species--such as abilities to explicitly conceptualize issues concerning edibility, draw logical inferences from the fact of a plant's being edible, have fully-fledged or explicit propositional attitudes such as belief, and so on. The following sections will show how this distinction of two levels of cognition might be articulated via some developments within the II framework.

6. Fodor, Concepts and Thinking About

The basic issue under discussion is that of how to distinguish, within an II-based

framework, low-level classificatory kinds of cognition from high-level, logically

structured propositional attitude kinds of cognition. One useful approach to the issue

proceeds via consideration of Fodor's well-known views on concept possession. Fodor

holds that possession of a concept of Fs, such as a concept of cats, is a matter of being

able to *think about* Fs (Fodor 1998, 2008). In particular, he denies that conceptual

possession is merely a matter of being able to classify or sort Fs from non-Fs, in a manner

that could be established by scientific tests of dispositions to behaviorally discriminate Fs

from non-Fs. I propose to defend Fodor on this issue from an unusual direction, as

follows.

Fodor himself typically poses the issue as closely related to the issue of conceptual

atomism versus various holistic forms of conceptual pragmatism that involve some form

of inferential role semantics for concepts. However, there is another potential strategy of

defense. This is to claim that concepts cannot be adequately understood, or genuinely

possessed, independently of *fully propositional contexts* in which thinkers could employ

them. Understood thus, Fodor's claim would be that legitimate concept possession is

closely associated with *propositional attitude abilities*--abilities to think about Fs in

propositional ways, to have F-related beliefs or desires that are expressible in

propositional form, and so on. In contrast, it is arguable that mere discriminative sorting

abilities with respect to Fs and non-Fs are not sufficient to establish any such

propositional attitude abilities. On this broadly propositional view of concept possession,

sorting abilities alone are not sufficient for a thinker to possess propositional attitude abilities.

In terms of the II framework, the distinction could be characterized as one between two different ways in which a cognitive state could be indexed by a proposition. These could be distinguished as *first-order indexability* (FOI) versus *second-order indexability* (SOI). In simple, non-propositional cases of classification dispositions, ITP-based dispositions could be correct or incorrect, and hence first-order indexable (FOI) by a true or false proposition. For example, if an organism can reliably identify and eat foods that are nutritious for it, and behaviorally discriminate those items from non-food items, to that extent its cognitive states are FOI.

However, the mere fact that an organism Z has FOI states does not entail that, in the full propositional-attitude sense of second-order indexability (SOI), Z itself actually possesses the relevant concepts that would feature in its having propositional attitudes of belief and desire etc toward relevant items. Or, in my imputed Fodorian terms, a mere ability by Z to provide evidence that it can behaviorally discriminate nutritious versus non-nutritious items does not, by itself, thereby provide any evidence that Z possesses explicit concepts of food, or of what it is to be nutritious or non-nutritious, nor that Z itself possesses any relevant propositionally based beliefs or desires--since the possession of such would require that Z did indeed possess the relevant concepts.

But exactly what is involved in an organism being capable of not merely FOI states, but also at least some SOI states? Fodor's representational theory of mind (RTM) involves a claim that genuine concept-possession is a metaphysical issue related to computational symbol-processing abilities (Fodor 1998, 2008). To genuinely possess the concept of CAT is to have a mental symbol "CAT"--a mental particular or token of the symbol type CAT--that one can employ in propositional symbolic computations that are causally related to cats. On this account, to *think about* cats, or a cat, is to employ a token of the mental symbol CAT in computations causally related to cats, or a cat. As for mere sorting or behavioral discrimination abilities, generally Fodor's attitude is to regard their invocation as part of a failed rival theory to his RTM--an alternative, pragmatic account of concept possession and thinking that, among other things, fails to adequately address issues of logical structure such as that of the compositionality of thoughts and concepts.

What is now needed is an alternative, II-based account of the difference between full propositional thought and basic sorting or classification. Potentially such an account would have an advantage over the Fodorian account, in that it could provide legitimate roles both for low-level classification dispositions and for high-level propositional thought--as complementary elements in cognition, rather than as incompatible factors in competing overall accounts of cognition.

Next, recall that, according to the II, perceptual causality may be explained in terms of the interactive theory of perception (ITP). The basic structure of the ITP itself, as a causal framework that involves both input causality and output causality factors, strongly

suggests the following potential account of the difference in cognitive levels. Simple, non-propositional classification cases would be cases close to a behavioral paradigm, in which only input and output factors would need to be considered, with no significant intervening cognitive factors. In terms of a computational model, these purely behavioral cases would typically involve a simple lookup table, in which a particular kind of input data configuration would automatically trigger a stereotypical output classification disposition. No propositional thought or belief would be involved in such cases, because they would be nothing more than cases of purely routine sorting-related behavior.

On the other hand, presumably higher-level propositional cases would be those in which *some significant kinds of cognitive processing* intervene between sensory inputs and behavioral outputs. Arguably it is at this specific point that Fodorian intuitions about the importance of concepts as abilities to *think about* objects falling under a concept become relevant. For example, an ability to think about cats, as opposed to an ability merely to be able to classify some items as cats or non-cats when they are perceived, presumably involves an ability to imagine a cat, or to conceive of what it would be for something to be a cat, *independently* of any current perceptions of worldly objects or classifications thereof.

A closely related point is as follows. In order to possess a general concept, such as that of a color or a natural kind, one must be able to understand *what it would be for an object to satisfy or not satisfy the concept* (Evans 1982, Peacocke 1992, Fodor 1998). For example, in order to possess the concept of redness--so that the concept is cognitively

salient for oneself, as one that one understands and can use in formulating truth-evaluable thoughts--one must be able to understand what it would be like for a typical object either to be, or not to be, red.

Returning now to perceptual cases, these initial considerations concerning concept-possession suggest the following high-level role for concepts in perception. Rather than input sensory data immediately triggering output classification dispositions, as in non-conceptual, low-level cases, instead the input data is subject to some specifically conceptualized kinds of cognitive processing. These might involve, for instance, a generate and test mechanism, in which the perceiver's conceptual abilities are use to generate putative or test instances falling under a concept, which putative instances could then be compared to the actual input data for matches.

A range of concepts could thus be tested for matching with the input data, and when a match is found, this presumably would cause appropriate, concept-specific output classification dispositions to be formed or activated. This would be one way in which a specifically cognitive, conceptual kind of processing could intervene between sensory inputs and classificatory outputs. The next section will suggest a further kind of intervening cognitive processing, specifically of a propositional-attitude kind.

7. Cognitive Processing of Propositional Attitudes

The approach to be adopted to propositional attitudes, such as that of believing a proposition, is one in which the cognitive structure of belief is explained as a further refinement of the intervening conceptual structure that mediates between sensory inputs and classificatory dispositional outputs. As just discussed above, possession of a concept C may be regarded as involving an ability to generate sample or putative instances of concept C, so that these may be compared with incoming sensory data to find a match. Belief could then be regarded as a further refinement, in which an additional assessment is made as to the *epistemic quality* of the incoming sensory data. For example, a distant object X, or one observed under poor lighting conditions, might trigger a conceptual match with a putative instance of concept C, but nevertheless *not* cause the perceiver to become disposed to classify object X as falling under concept C, because of the poor quality of the available sensory evidence.

On this account, the concept of belief plays a *normative epistemic* role in cognition. The difference between perceptually caused low-level, unconceptualized classification dispositions, and high-level, perceptually caused beliefs that may also eventuate in similar classification dispositions, is that the beliefs involve an intervening assessment of the quality of the evidence provided by the sensory inputs. On this approach, an epistemically responsible perceiver is one who becomes disposed to classify an object X as being F only if the available evidence, including sensory evidence, provides adequate grounds for his believing the proposition that X is F. Or, in terms of the II theory, the difference between unconceptualized first-order indexing (FOI) and fully conceptualized, propositional-attitude-involving second-order indexing (SOI) may be explained in the

same way. In this manner, the II theory can be appropriately integrated with standard epistemic issues concerning naturalistically justified belief (Sosa and Kim 2000, Steup 2008).

On this II account, a unified account of perceptual representation is still available, as always involving classification dispositions that are indexed by appropriate propositions. The higher-level additions provided by SOI, over and above FOI factors, affect only the mediating conceptual and propositional attitude factors, which modify the conditions under which a perceiver acquires, or does not acquire, relevant classification dispositions in the case of a given sensory input.

This epistemic approach to propositional belief may gain some intuitive support from the following two-factor propositional model, which invokes a concept of *epistemic seeming* *(*Chisholm 1957, Dilworth 2006b). Suppose that a perceiver P possesses a concept C that can be used in generating putative examples of Cs. Then arguably incoming perceptual data could *epistemically seem* to P to be a case of an object satisfying C, in that P could generate a putative test example of C, and find that the incoming perceptual data matches that putative example. However, this kind of initial, specifically conceptual cognitive processing of an item X--as seeming to be C--does not by itself amount to a *belief* that X is C. For on the current account, belief requires also that the putative or seeming fact of X being C should be taken by the perceiver to provide reliable evidence that X actually is C.

Consequently, we must distinguish a mere perceptual seeming--of object X seeming to be

C--from a full belief that X is C, as follows. In order for perceiver P to believe that X is

C, he must believe that the apparent or seeming C-hood of X provides reliable evidence

of--or corresponds with--how X actually is. So both the seeming or putative fact of X

being C, and whatever is the actual fact concerning the possible C-hood of X, must play a

role in perceiver P's belief, since in normal cases it is primarily X's perceptually seeming

to be C that provides P's evidential basis for his belief that X is C.

These points may be summarized as follows. The difference between a non-propositional

case of X seeming to be C to a perceiver--i.e., a seeming or putative fact of X being C--

and a perceptually based, fully propositional belief that X is C, is that in the latter case

the perceiver must believe that the putative fact (PF) of X being C is reliable evidence of,

or corresponds with, an actual fact (AF) of X being C.   So the non-propositional seeming

or putative fact PF of X being C is only one component in the fully propositional belief

that the PF corresponds with an AF of X being C. Also, as one would expect, behavioral

evidence that a perceiver holds such a belief would typically be provided by his acquiring

C-related classification dispositions with respect to X, when X perceptually seems to him

to be C.

On the other hand, if such classification dispositions were not acquired under those

conditions, this would provide evidence that perceiver P does not believe that X is C--in

spite of its seeming to him that X is C. Such 'gappy' cases--of perceptual seeming

without believing--would also provide a convenient explanation of known perceptual

illusions (Dilworth 2006b, Crane 2008). For example, in the Muller-Lyer illusion of
equal lines between differing arrowheads--"<--->' versus ">---<" the lines may continue
to perceptually seem to be of unequal lengths, even though the knowledgeable perceiver
does not believe this to be the case. The current two-factor, PF/AF propositional model
provides a useful model of the difference between non-propositional, fact-like perceptual
seeming on the one hand, and two-factor, fully propositional believing on the other hand.

Another useful aspect of the two-factor PF/AF propositional model is that the behavioral
evidence for perceptual seeming can in principle be distinguished from the behavioral
evidence for full propositional believing--since the two kinds of evidence can come apart,
as just noted, in some perceptual illusion cases. Also, there is a clear cognitive
distinction between input data caused by objects, such as the raw initial data provided by
the sensory organs of a perceiver, and the eventual behavioral output data that provides
evidence of acquisition of relevant dispositions toward the relevant objects.

So, to summarize the overall evidential situation, the II theory as augmented by the two-
factor propositional model employed enables us to distinguish three kinds of cases. 1)
First-order indexing (FOI) cases in which unconceptualized, C-related classification
dispositions of an object X result from perception of X; 2) Second-order indexing (SOI)
cases in which both conceptualized seeming and full propositional believing occur, in
that C-related classifications do result from the perceptual seeming; and 3) degenerate or
mixed cases, in which conceptualized seeming--of X seeming to be C--occurs without
believing, in that no C-related classification dispositions of X result. However, for

26

theoretical simplicity, it might be arguable that cases in this third category could usually

be assimilated to second-category belief cases in which X seeming to be C causes a belief

that X is *not* C--i.e., causes dispositions to classify X as being non-C rather than as C.

8. The Augmented II Theory and the Representational Theory of Mind

From now on, the II theory, as supplemented by the two-factor PF/AF model of

propositional attitudes, will simply be referred to as the *augmented* II (AII) theory. This

section briefly considers the resulting overall relations of the AII theory to a Fodor-style

representational theory of mind (RTM), as initially discussed in section 5. To begin, the

last two sections have established that the AII could fully support a broadly Fodorian, or

at least Fodor-inspired, distinction between low-level sorting dispositions on the one

hand, and full-fledged propositional attitudes on the other hand, which deploy explicit

concepts that can be used in thinking about objects.

As for the more specific details of Fodor's proposals, such as his claim that genuine

possession of the concept of CAT involves employment of a token of CAT in symbolic

computations that are causally related to cats, the AII theory could remain neutral or

agnostic on the issue. This is so because the AII theory is a purely functional theory of

cognition that is independent of specific implementation details, such as that of whether it

is implemented in the form of computation over explicit concept symbols, as in the

Fodorian view, or in some other manner, such as some non-symbolic neural-net

implementation. Arguably this gives the AII theory a distinct advantage over a Fodorian

RTM, in that the AII-style account of concept deployment in propositional attitudes is

also consistent with various competing, non-RTM views of cognition (Wilson and Keil

2001, Thagard 2008).

The AII theory is also closely aligned with a Fodor-style RTM on intentionality issues, in

spite of the official intentional eliminativism of the basic II theory. For, though the II

theory denies there are any concrete intentional states, it nevertheless re-locates or re-

identifies intentional relations as legitimate and integral features of purely abstract

propositional structures that index relevant cognitive states. Consequently, Fodor's

industrial-strength intentional realism (Fodor 1990, etc.) can be closely paralleled by the

II's metatheoretical scientific realism (Dilworth 2008, 2009), in which abstract

propositions are viewed as being part of the realistic, metatheoretical presuppositions of

any adequate scientific theorizing, including theorizing about cognition. As a result, both

the RTM and the AII theory closely identify cognitive abilities with abilities to think

about particular objects, and ascribe specific properties to them in truth-evaluable ways

that can be expressed in objective singular propositions.

Another close alignment issue is that of conceptual atomism versus holism. This

alignment may be unexpected, because, as noted previously, Fodor regularly argues or

assumes that dispositional accounts of cognition inevitably lead to conceptual holism

(Fodor 1998, 2008). However, the following preliminary points are sufficient to

demonstrate the potential close alignment. In a broader context, the II theory would be

situated in an account of the evolutionary benefits of kinds of perceptual and cognitive processing that are *propositionally indexable*. Indeed, cognitive science is mainly a study of those indexable kinds, and arguably any study of logically acceptable, compositionally related concepts would necessarily--for broadly Fodorian reasons (Fodor and Lepore 2002)--be exclusively confined to the study of indexable cognitive states, whose contents could be accurately described or indexed in propositional terms.

Consequently, in order to participate in the evolutionary benefits of indexable cognition, species--such as us humans--must develop, or have developed, kinds of ITP-based dispositional causality that are indeed indexable, whether in low-level FOI or high-level SOI forms. But, in order for concrete cognitive states to be propositionally indexable at all, they must be unobjectionably conceptually atomistic to the same extent as are the atomistic concepts that are involved in the abstract propositions that index those concrete states. As a result, arguably the AII theory is just as strongly committed to conceptual atomism as is Fodor himself, in spite of the underlying differences between a Fodorian RTM and the dispositionally based AII theory.

9. The Dispensability of Propositional Indexing

The II theory was developed via an analogy between abstract mathematical indexing of quantitative physical states, and abstract propositional indexing of truth-evaluable cognitive states (Dilworth, 2008a, 2009). This section will further investigate the

comparison with respect to indispensability issues of the following kind. It is widely assumed that some kinds of mathematical indexing are *indispensable* in a naturalistic worldview, in that they index or abstractly categorize important scientific properties that could not be described or conceptualized at all without the aid of mathematical concepts (Quine 1980, Colyvan 2008). On this indispensability view, any quantitative scientific concepts of absolute or relative amounts of matter could only be described at all with the aid of numeric indexing. If one object A weighs more than another B, then the heavier item A must have a *greater quantity* of units of matter than the lighter item B--i.e., A is indexable by a higher numeric quantity of units of matter than is B. Thus the indispensability assumption is that the quantitative relative difference in masses must be described via the employment of abstract numeric concepts, even though--strictly speaking--the physical objects could not actually instantiate numeric concepts.

If the general theoretical presuppositions of the II theory are correct, a similar indispensability thesis holds for the direct relations of propositions and propositional concepts to the world, as well as for numeric concepts. The II theory claims that the only genuine kinds of representation and intentionality are purely abstract kinds. But since we cannot describe or think about the world, according to the II theory, without being in concrete cognitive states that are indexed by such abstract propositions, those propositions too are semantically indispensable to our best scientific accounts of the nature of the physical world.

However, the primary focus of the II theory is not on the indispensability of abstract propositions themselves in directly representing, or directly being about, the world. Instead the theory is a theory of how truth-evaluable cognitive states are *indirectly* indexed by such propositions that are directly about the world. So for clarity, the concept of propositional indexing, or simply of indexing, will henceforth be restricted to this primary kind of propositionally mediated, *indirect* indexing or correlation between a cognitive state and the world. Then the main topic for this section is whether those concrete cognitive states that are thus indirectly propositionally indexed can themselves be completely described without invoking or presupposing any content-related indexing concepts.

Now it might mistakenly be thought that the basic theoretical stance of the II theory itself --namely, that purely causal issues about cognitive states must be sharply distinguished from issues about the abstract propositions that index them--has already prejudged this issue in the negative. On such a mistaken view, indexing descriptions of concrete cognitive states would have to be dispensable, in order to ensure that cognitive states could themselves be purely physical, non-intentional states.

But such a view is indeed mistaken, as is shown by the case of direct mathematical indexing or correlation of concrete quantitative states with abstract numerical sequences. Even though arguably descriptions of such direct mathematical correlations are indispensable to descriptions of physical reality, their being so has no bearing on the purely physical status of the physical phenomena themselves that they abstractly

describe. And similarly, it is, at least initially, an open theoretical possibility that purely physical cognitive states that are also propositionally indexable may have physical aspects for which abstract indexing descriptions are indispensable.

Nevertheless, in spite of this open theoretical possibility of an indispensability claim for indexing of cognitive states, there are some compelling, broadly methodological issues that should make us prefer--at least initially--an alternative claim of the *dispensability* of indexing concepts in descriptions of cognitive states. One main methodological issue is that of naturalistic uniformity. Since physical objects and states in general can be fully described without invoking any propositional indexing concepts, it would be suspiciously ad hoc to claim that one specific group of physical states--namely, indexable cognitive states--happens also to involve indispensable indexing in descriptions of its members.

Consequently, if the II theory is to maintain its substantive naturalistic credentials, it should, at least initially, treat concrete cognitive states as being fully on a par with any other physical states, and so as involving, if at all, only dispensable indexing concepts. To be sure, it cannot be ruled out in advance that this policy might break down at some point. For example, there might eventually be some plausible arguments to the effect that through evolution, some cognitive states have acquired complex, emergent physical properties whose descriptions happen to indispensably involve indexing concepts. But until such arguments become available, this possibility can reasonably be ignored for the present.

The acceptance of this *indexing despensability* (ID) thesis for cognitive states has some important implications for various issues, including for issues of intentional realism with respect to cognitive states. Previous sections have shown that the II and augmented II (AII) theories are broadly comparable in various respects to a Fodor-style intentional realism. For example, on the II account, intentionality is real (though abstract), and there is in general a fact of the matter as to what is the propositional content of a given concrete belief state. However, we must now distinguish *strong* intentional realism for cognitive states--involving a claim that intentional descriptions of such states are indispensable-- from the currently adopted ID view, which could be described as a species of *weak* intentional realism.

For weak intentional realists, intentional descriptions of appropriate cognitive states are possible and legitimate, but nevertheless they are *theoretically optional* or *dispensable*-- nothing would be left out of a complete physical description of those states, even if no indirect, indexing-style or content-based intentional concepts were included in those descriptions. By contrast, on a Fodor-style strong intentional realist view--according to which intentionality is physically real, via some appropriate causal reduction that preserves the relevant intentional structures--the relevant cognitive states necessarily have intentional properties, so that intentional descriptions of those physical states are indispensable.

As defined here, dispensability and indispensability claims are ontological or metaphysical claims, pertaining to the actual nature of the physical states that are being

described. However, such specifically ontological issues must be distinguished from

broader scientific issues of *explanatory* dispensability or indispensability for various

kinds of scientific purpose. For example, the II theory claims that an explanation of what

it is for a cognitive state to be semantically meaningful indispensably involves both

concrete cognitive states and their indirect indexing by abstract, intentionally based

propositions.   But such an explanatory theory of meaningful cognition involves, of

course, much more than simply a basic inventory of the physical nature of the relevant

cognitive states. So it is legitimate to hold that propositional indexing may be

explanatorily indispensable, in spite of its being *ontologically* dispensable in basic

physical descriptions of cognitive states.

Perhaps the closest analogy in the extant literature for the current defense of weak

intentional realism is Dennett's well-known view of intentionality as a practically

indispensable but theoretically eliminable stance toward cognitive states--i.e., his doctrine

that intentionality is an *intentional stance* (Dennett 1987). Dennett's view remains

controversial, in that any number of commentators have been baffled in trying to pin

down exactly what his view is and what it entails (Dahlbom 1995, Ross and Brook and

Thompson 2000). The II theory could help to clarify this issue, in that it occupies a clear

intermediate position between a strong intentional realism and a pure instrumentalism

about cognitive intentionality--both of which extreme positions Dennett rejects.   The II

theory denies strong intentional realism about intentionality in denying that cognitive

systems themselves have any indispensable intentional properties, even of an abstract

kind. Nevertheless, as with Dennett's view, its indexing approach is not a pure

instrumentalism either, in that it claims that there is a substantive, experimentally based fact of the matter as to which cognitive states are indexed by which propositions.

It is in this specific respect that the II theory can improve on and clarify what is correct in Dennett's denial of instrumentalism. According to the II theory, the cognitive benefits of using intentional concepts are securely grounded in the scientific benefits of being able to describe the results of experiments on dispositionally based perceptual abilities in propositional terms that specify *what a subject perceives or believes*--hence furthering the overall predictive and explanatory goals of science considered as an abstract theoretical structure of propositions.

 Nevertheless, this kind of intentionally based scientific indexing remains dispensable, in that a complete scientific account of the relevant cognitive abilities could instead be given in terms of propositions specifying no more than the purely dispositional, four-factor causal structures addressed by the interactive theory of perception (Dilworth 2004, 2005a, b, c, 2006a, 2008a, 2009, forthcoming), or similar purely causal theories of cognition.

10. The Naturalistic Status of Classification Dispositions

To conclude this paper, a potential criticism should be addressed, namely that there seem to be descriptions of cognitive dispositions relevant to intentional indexing that do

indispensably describe those dispositions in propositional terms. If that were so, there would be perceptual or other cognitive dispositions involving intentional properties that could not be identified independently of propositional concepts. For example, some committed naturalists might object to the explanation in Dilworth (2008, 2009) of how *classification dispositions* may be propositionally indexed, on the ground that the concept of classification is itself an irreducibly intentional concept (e.g., Fodor 1998). Also, the criticism might continue, the II theory itself compounds the problem via its invocation of a concept of propositional indexing. In this compounded form, the objection might be that an appeal to classification dispositions is just a more covert form of appeal to--for example--a disposition of scientists to *indirectly propositionally index* the intentional perceptual dispositions of their experimental subjects. But, the objection would continue, such dispositions to propositionally index intentional states are clearly dispositions having indispensable intentional propositional properties. Hence, the objection would conclude, intentional indexing of cognitive states is indispensable rather than dispensable, and consequently intentional properties of cognitive states cannot be eliminated after all.

This objection is a useful one because it illustrates one kind of explanatory role of intentional concepts. It would typically be intuitively explanatory and helpful to describe some dispositions as classification dispositions, or to explicitly identify a disposition of scientists to *indirectly propositionally index* the intentionally relevant perceptual dispositions of their experimental subjects, whether or not the objection is well-founded. As it happens, the objection is not well-founded, because it ignores the fact that the basic concept of a causal disposition is a purely extensional concept that is *description-*

36

independent. It does not matter how a disposition, or its activation conditions, are described, as long as a causal disposition is clearly identified that would be causally manifested under appropriate conditions.

For example, a disposition to put a red object in a box with other red objects could be described in various ways, including describing it as a manifestation of a classification disposition. But the disposition and its manifestation conditions can be identified and scientifically studied independently of any particular way of describing it. Consequently, *the disposition itself* does not have any indispensable intentional properties, even if some ways of describing it do characterize it in intentional terms in order to reap the explanatory benefits of so doing.

To be sure, the issue of whether the concept of classification is indeed irreducibly intentional, as assumed by the objection above, is a complex one that would require much more discussion to resolve.   But since dispositions are description-independent in any case, it would make no difference to the II theory or its foundations how that discussion might turn out.

A more technical issue concerning the status of dispositions should also briefly be addressed. Some have held that dispositions, as with some other basic causal properties or constructs, must be intrinsically defined (Lewis 1997, Molnar 1999). On such a view, strictly speaking there could not be genuine or basic dispositions involving extrinsic relations to other objects, as postulated by the interactive theory of perception (ITP).

However, there are now enough countervailing arguments in the literature (McKitrick 2003, Handfield 2009) to render that objection moot.

Also, even if the intrinsicalist view of dispositions were eventually to prevail, a fallback position of the following kind would be available to II theory supporters. The interactive causal structures linking a perceived object X to its perceiver Z and her X-related causal relations would be in good scientific standing in any case, even if, strictly speaking, there could not be any X-related causal *dispositions* as such. Under those conditions, some more complex causal constructs might need to be invoked to fully explain the legitimate naturalistic basis of the ITP. Nevertheless, the concept of an extrinsically defined disposition could still persist as a useful explanatory simplification of that more complex, fully naturalistic account. It should not be forgotten that the II theory, along with its augmented (AII) version, is an applied, naturalistic semantic theory, not a basic metaphysical treatise concerning the nature of scientifically fundamental dispositions. So here too, the legitimate status of the AII theory as a semantic theory is not dependent on, or potentially impugned by, whatever might be the resolution of such more technical disputes.

References

Chisholm, R. (1957). *Perceiving.* Ithaca, N.Y.: Cornell University Press.

Colyvan, M. (2008). Indispensability Arguments in the Philosophy of Mathematics. *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2008/entries/mathphil-indis/>.

Crane, T. (2008) The Problem of Perception. In *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2008/entries/perception-problem/>.

Cummins, R. 1996. *Representations, Targets and Attitudes*. Cambridge, MA: MIT Press.

Dahlbom, B. (ed.) (1995). *Dennett and his Critics.* Oxford: Blackwell.

Dennett, D. (1987). *The Intentional Stance.* Cambridge, MA: MIT Press.

Dilworth, J. (2004). Naturalized Perception Without Information. *The Journal of Mind and Behavior,* 25, 349-368.

Dilworth, J. (2005a). The Reflexive Theory of Perception. *Behavior and Philosophy, 33,* 17-40.

Dilworth, J. (2005b). Perceptual Causality Problems Reflexively Resolved. *Acta Analytica,* 20, 11-31.

Dilworth, J. (2005c). A Naturalistic, Reflexive Dispositional Approach to Perception. *The Southern Journal of Philosophy,* 43, 583-601.

Dilworth, J. (2006a). A Reflexive Dispositional Analysis of Mechanistic Perception. *Minds and Machines,* 16, 479 - 493.

Dilworth, J. (2006b). Perception, Introspection and Functional Consonance. *Theoria*, 4, 299-318.

Dilworth, J. (2008a). Semantic Naturalization via Interactive Perceptual Causality. *Minds and Machines,* 18, 527-546.

Dilworth, J. (2008b). The Propositional Challenge to Aesthetics. *British Journal of Aesthetics,* 48, 115-144.

Dilworth, J. (2008c). The Abstractness of Artworks and its Implications for Aesthetics. *The Journal of Aesthetics and Art Criticism*, 66, 341-353.

Dilworth, J. (2009). Semantics Naturalized: Propositional Indexing Plus Interactive Perception. *Language & Communication,* 29, 1–25.

Dilworth, J. (Forthcoming). Realistic Virtual Reality and Perception. *Philosophical Psychology.*

Dretske, F. (1981). *Knowledge and the Flow of Information.* Cambridge, MA: MIT Press.

Evans, G. (1982). *The Varieties of Reference.* Oxford: Clarendon Press.

Fodor, J. (1987). *Psychosemantics.* Cambridge, MA: MIT Press.

Fodor, J. (1990). *A Theory of Content and Other Essays.* Cambridge, MA: MIT Press.

Fodor, J. (1998). *Concepts: Where Cognitive Science Went Wrong.* Oxford: Clarendon Press.

Fodor, J. (2008). *Lot 2: The Language of Thought Revisited.* Oxford: Clarendon Press.

Fodor, J. and Lepore, E. (2002). *The Compositionality Papers.* Oxford: Clarendon Press.

Handfield, T. (ed.) (2009). *Dispositions and Causes.* Oxford: Oxford University Press.

Lewis, D. (1995). Finkish Dispositions. *Philosophical Quarterly,* 47, 143-158.

McKitrick, J. (2003). A Case for Extrinsic Dispositions. *Australasian Journal of Philosophy,* 81, 155-174.

Molnar, G. (1999). Are Dispositions Reducible? *Philosophical Quarterly,* 49, 1–17.

Peacocke, C. (1992). *A Study of Concepts.* Cambridge, MA: MIT Press.

Quine, W. (1980). On What There Is. Reprinted in *From a Logical Point of View*, 2nd edition, Cambridge, MA: Harvard University Press, pp. 1–19.

Ramsey, W. (2007). *Representation Reconsidered.* Cambridge: Cambridge University Press.

Ross, D. and Brook, A. and Thompson, D. (2000). *Dennett's Philosophy: A Comprehensive Assessment.* Cambridge, MA: MIT Press.

Sosa, E. and Kim, J. (2000). *Epistemology: An Anthology.* Oxford: Blackwell.

Steup, M. (2008). The Analysis of Knowledge. In *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2008/entries/knowledge-analysis/>.

Thagard, P. (2008). Cognitive Science. In *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2008/entries/cognitive-science/>.

Wilson, R. and Keil, F. (2001). *The MIT Encyclopedia of the Cognitive Sciences.* Cambridge, MA: MIT Press.