The Journal of Mind and Behavior
Autumn 2004, Volume 25, Number 4
Pages 349–368
ISSN 0271–0137

Naturalized Perception Without Information

John Dilworth

Western Michigan University

The outlines of a novel, fully naturalistic theory of perception are provided, that can explain perception of an object X by organism Z in terms of *reflexive causality*. On the reflexive view proposed, organism Z perceives object or property X just in case X causes Z to acquire causal dispositions *reflexively directed back upon* X *itself*. This broadly functionalist theory is potentially capable of explaining both perceptual representation and perceptual content in purely causal terms, making no use of informational concepts. However, such a reflexive, naturalistic causal theory must compete with well entrenched, supposedly equally naturalistic theories of perception that are based on some concept of *information*, so the paper also includes some basic logical, naturalistic and explanatory criticisms of such informational views.

Keywords: perception, naturalism, information

One main purpose of this paper is to present the outlines of a novel, fully naturalistic theory of perception that can explain perception of an object or property X by organism Z in terms of *reflexive causality*. On the reflexive view to be proposed, organism Z perceives object X just in case X causes Z to acquire causal dispositions reflexively directed back upon X itself. Thus the theory is a broadly functionalist theory, which gives a theoretical role to all three factors of causal inputs, dispositional cognitive intermediaries and behavioral outputs. The theory is outlined below, and it is potentially capable of explaining both perceptual representation and perceptual content in purely causal terms that make no use of informational concepts.

However, such a naturalistic causal theory must compete with well entrenched, supposedly equally naturalistic theories of perception that are based on some concept of *information*, whether explained in terms of nomic

My thanks to the editor and reviewers of the *Journal of Mind and Behavior* for very helpful comments. Requests for reprints should be sent to John Dilworth, Ph.D., Department of Philosophy, Western Michigan University, Kalamazoo, Michigan 49008. Email: dilworth@wmich.edu

covariation, as with writers such as Dretske (1981) and Fodor (1990), or using a weaker statistical (probability of less than 1) concept of information, such as in Millikan's recent proposal based on the concept of a "local natural sign" (Millikan, 2001, 2004).[1] As applied to perception, such informational views claim that the representational contents of perception are entirely to be explained in terms of such sources of information, items of which are by definition true or correct.

Given the entrenched status of such informational views, no alternative naturalistic view can hope to establish itself purely on its own merits. Some vigorous criticism, including pointed reminders of basic flaws in those informational views, is also necessary. Hence some relatively neglected, but nevertheless arguably conclusive kinds of criticism, will be supplied.

Informational views have also been criticized on additional grounds, a prime complaint being that they make no allowance for the possibility of perceptual *misrepresentation* — or more generally, that both perceptual and cognitive contents have a genuinely semantic status, such that they may be either true or false (or correct or incorrect) with respect to the objects or states of affairs that they are about, so that at best necessarily correct informational content could only play a subsidiary role in perception or cognition (Cummins, 1989, 1996). For example, Cummins (1996) claims that such theories conflate the content of a representation with the worldly situation that it represents, hence making misrepresentation impossible, while also claiming that purely indicative informational content must be sharply distinguished from the genuinely representational content of perception (Cummins and Poirier, 2004). However, though such misrepresentation-related objections are indeed significant, arguably the logical, naturalistic and epistemic reasons for rejecting informational approaches to be presented here are at least as fundamental.

As for the rest of the paper, later sections show how the reflexive view could be extended so as to explain the *intentionality* of perception, including in pictorial cases, as well as potentially to phenomenal or conscious aspects of perception also, though the main concentration is on non-conscious factors. Thus broader issues about naturalistic accounts of conscious experience are not here addressed.

Logical, Naturalistic and Scientific Relevance Objections to Information

As a preliminary, the legitimate metaphysical status, and even the coherence, of the relevant concept of worldly semantic information may be questioned,

[1]Millikan (2001) describes information as "informationC," or correlational information, that necessarily is correct. See also the more comprehensive account in Millikan (2004).

so that it may be inapplicable in all possible worlds, and hence our own as well. Presumably there is a *fact* that object X is F, if the *proposition* "X is F" is true, and that proposition is also *true* if it is *known* to be true. But the concept of an always correct, yet worldly informational item "X is F," looks suspiciously like an incoherent conflation of these distinct metaphysical, semantic and epistemic concepts. Items of information are supposed to be "in" the world like facts, but also not be in, but instead *about*, the world like propositions — which can be true or false. Yet they are also viewed as being *necessarily correct*, and hence never false or incorrect, which looks like a confused analysis of what it is to know a proposition to be true.

Thus it is somewhat surprising that a concept which is, to all appearances, an unresolved mixture of such disparate logical characteristics, could have became respectable enough to be associated with supposedly naturalistic programs in biology and cognitive science. One main source of the assumed legitimacy of informational concepts is likely their confusion with standard semantic and epistemic concepts such as that of reliable true beliefs — that might have been false — which organisms are capable of acquiring through completely naturalistic, causal means. Another important source is likely a conflation of informational concepts with purely naturalistic concepts such as those of causal covariation or statistical correlation, which presumably have some legitimate role in perceptual theories (Dretske, 1981; Millikan, 2004), or in reliabilist theories of justification of belief (Goldman, 1976; Kornblith, 2002).

More broadly, informational concepts are arguably just the latest naturalistic attempt to explain perceptual contents purely in terms of *causal inputs* — specifically, incoming sensory data — as in traditional psychological and philosophical views of sensualism or sensationalism (Agassi, 1966). A broadly functional or behavioral view of perception, such as that to be presented, will reject such purely input-based theories in favor of a more comprehensive functionalist view that gives a proper place to all three factors of causal inputs, cognitive intermediaries, and behavioral outputs. Informational views constitute a particularly extreme form of purely input-based approaches to perception because of their additional logical flaws, but many of their defects may be found in any purely input-based theories.

For the sake of argument it will be assumed that the relevant concept of information is at least minimally coherent enough so that arguments as to its specifically empirical and scientific, rather than merely logical, failings may be presented. Also, though some of the criticisms to be presented would apply equally well to covariational or statistical accounts of perception, such purely naturalistic issues will have to be discussed elsewhere.

To begin, the initial claim to be denied is that the natural world *itself* contains or produces any information, or natural signs or representations having an information-based content, of a sort that could be causally acquired by

perceivers so as to internalize that same information about the world. Thus the initial thesis to be argued is that there is *no worldly information*. Both naturalistic and scientific relevance or plausibility arguments will be used to defend this thesis.

As an initial source of arguments in favor of this no-worldly-information view, consider a world containing no organisms capable of perception at all. In such a world, the scientific postulation of worldly information or informational content would be completely idle, or explanatorily irrelevant. At the same time, such a postulation would, as always, violate physicalistic or naturalistic standards of explanation, since there is no clear way in which to explain or reduce information or informational content, such as a supposed item of information that an object X has property F, to acceptable, purely physicalistic causal factors.

Next, consider a later stage of that same world, into which a perceiving organism Z has been introduced. Since the world itself has not changed, it could not be correct to argue that the mere addition of organism Z has somehow transformed the world into one including or producing worldly information. In particular, any *peripheral causal effects* of the world upon organism Z, such as a state S of retinal stimulation of its eyes, also could not somehow become information transmitting or producing in those eyes as such, since qualitatively the very same causal process could have existed in the world — namely, a stimulation of a retinal surface — prior to its becoming the case that a thus-stimulated retinal surface S exists that is also a peripheral part of a perceiving organism Z. Or in other words, the necessary objectivity and perceiver-independence of information assumed by its proponents entails that worldly information cannot suddenly appear out of nowhere just because a perceiver has become related to that world. Hence, in such a case, reasonable standards of localized explanation would require that any *information acquisition by organism* Z be explained, if at all, purely in terms of its own internal, non-peripheral workings. Thus any information now in the world must be localized within the internal activities and structure of organism Z.

Now consider a third stage of world A, in which a range of perceiving organisms under evolutionary adaptive contingencies are included, or have developed (this world might be our actual world). Again, such contingencies might, let us suppose, affect the perceptual information and content acquired by members of each species, but still the *world itself* could not have acquired any informational properties, which, if there are any at all, must continue to be localized within the relevant organisms.

The overall argument now shifts to a consideration of the postulated localized *internal* information — information somehow produced by purely causal, non-informational means *within* an organism, as discussed above. But informational concepts are just as naturalistically unacceptable when applied to

internal processing as they were to external worldly causal events. Further, the explanatory relevance of postulating specifically internal information at this stage must also be questioned, since the postulation of organisms internally acquiring information only makes sense if such information could be transmitted to the organism *from some external source* — but such external transmission has already been ruled out. Thus external (worldly) and internal information stand or fall together, and, as shown, both for naturalistic and explanatory relevance reasons both must fail.

The argument used here is related to, but distinct from, extant arguments that information would be epiphenomenal or causally inert (Dretske, 2000), and hence of no explanatory value for that reason (on which see the next section). Instead the argument used is a recursive one, focused on the conditions under which introduction of a concept of information could be *explanatorily justified.* In the "zero" case, with no perceiving organisms, it could not be justified. In the next stage, with a single perceiving organism, it equally could not be justified — and so on for arbitrary increases in the number and complexity of organisms. Hence it cannot be justified at all.

Further arguments against the scientific relevance or plausibility of information concepts will be given below, interspersed with the suggested purely causal perceptual alternative. The overall claim will be that not only philosophical naturalists, but also scientists generally, should reject the relevant informational concept, and any perceptual theories that use it. Arguably naturalistic programs which postulate worldly information in some form or other, such as the "informational semantics" of Dretske, Fodor et al., or Millikan's adaptive view, are naturalistically suspect because of their use of such informational concepts.

The Irrelevance of Information to Perceptual Evidence

This brief section will provide an additional, specifically epistemic reason as to why the postulation of information is explanatorily irrelevant to a scientific theory of perception. The argument will be that the kinds of evidence that are actually available for the support of a perceptual claim that organism Z has perceived object X never require the postulation of information acquisition or processing by Z.

The basic epistemic claim, to be supported by the rest of the paper plus general naturalistic considerations, is that all third person, scientifically accessible perceptual *evidence* — i.e., evidence that perception has actually occurred — is broadly causal or behavioral in nature.[2] For example, if organ-

[2]Hence ignoring first person claims of immediate, self-evident and indubitable knowledge that one is currently perceiving X — which claims are in any case readily falsifiable by perceptual illusion experiments.

ism Z turns its head toward a predator X that is close to Z, but subsequently Z engages in no predator-related behavior whatsoever — such as "freezing" or fleeing — then a claim that Z *perceived* the predator Z has been defeated (assuming of course normal conditions, such that the rest of its body has not been paralyzed, and so on). On the other hand, if Z does freeze or flee immediately subsequent to turning its head toward predator Z, that behavioral evidence is sufficient to establish a claim that Z did perceive X (again, assuming normal conditions).

Consider each of those cases — no predator-related behavior, versus positive predator-related behavior — with respect to informational considerations. In either case, it makes no *perceptual* difference whether Z did, or did not, acquire information about predator X immediately subsequent to its head-turning — assuming for the sake of argument that there is such a thing as information-acquisition — since without necessary behavioral evidence, such information gain is not sufficient for perception, and with sufficient behavioral evidence, it would not be necessary for the occurrence of perception. Hence information acquisition is purely epiphenomenal, having no place in the causal order of nature, and thus it is scientifically irrelevant to perception, even if it were physically or naturalistically admissible.[3]

Perception and Reflexive Causality

If all evidence of perception is causal or behavioral evidence, as proposed above, then the most appropriate kind of theory of the nature of perception would be one that analyzed perception in purely naturalistic, causal terms, using nothing more in its ontology than particular causal relations between a perceiver Z and an object X, plus any purely causal dispositions of Z that were themselves causally associated with the perceptual situation. Using only such causal materials, there would be no room for views of perception as fundamentally based on the acquisition of information, conscious experience, irreducible intentional perceptual content, and so on. Instead a straightforward dispositional view of perception would have to be developed, in which what it is for organism Z to perceive object X is for Z to be caused by X to change its causal dispositions toward X in some way. I shall argue that a plausible theory of this kind is indeed available.

However, as a preliminary, the case for such a view can be strengthened if there are arguments available, independent of naturalistic considerations, as to why these kinds of causality should be taken to be involved in the very

[3]Dretske (2000) attempts to argue that information at least is explanatorily relevant in cases of genuine learning, but his arguments do not address the current more global objections to information.

analysis of the concept of perception.[4] One such condition is familiar, in that it is generally agreed that it is at least a necessary condition for a person or organism Z to perceive an object X is that X *causes* the relevant perceptual state S in Z. But of course there are plenty of states of Z, caused by X, that do not count as perceptual states, let alone as perceptions of X itself. So something more than adequate causal inputs must be required to explain perception.

Independent of causal considerations, it would generally be accepted as a minimal intuitive addition that it is a particular *representational* state S, which represents, or is *about*, the perceived object X, which must be caused by X. So the question is whether this additional factor of representation of X, or of being about X, generates in turn any additional necessary *causal* factor in the analysis of perception of object X. I shall claim that it does, in that the only objective, third person criteria or kinds of evidence for the presence of X-related intentional or representational aboutness in Z's state S are themselves causal, and provided by *X-related* causal dispositions in Z, that are in turn themselves caused by X — so that a kind of *reflexive* causality is necessarily involved in perception.

The intuitive idea is that genuine perception of X by Z involves Z being caused to have causal dispositions toward X itself, the activation of which would involve Z in turn causing some change in X (in the simplest case). With respect to evidence of perception, arguably the only conclusive evidence that animal Z has perceived food item X is if Z attempts to do things such as to *directly causally interact* with X in some way, such as by eating the food X, or hiding it for later use, and so on. On this account genuine perception involves not just causal inputs, but also modified internal states — dispositional changes — plus appropriate, X-related causal outputs as well.

This account needs to be immediately supplemented by "negative" X-related causal dispositions, for example a disposition to *refrain* from causal interaction with X when that interaction would otherwise occur, such as if Z is about to collide with object X, and its perception of X consists in its being caused by X to acquire a disposition, immediately activated, to avoid colliding with X. Here again, to supplement and complete the previous account, the best possible, and arguably the only conclusive, evidence that animal Z has perceived object X is if Z either attempts to avoid X, or attempts to interact with X.

The issue of naturalism may now be raised again. A purely naturalist theory of perception would claim that the two relevant causality factors involved in the analysis of perception, namely that X causes perceptual state S in Z, and that state S consists of X-related causal dispositions, are all that is needed to explain the representational, X-related aboutness of perception. Such a theory will be defended here. Thus the main innovation in the pre-

[4]Lowe (2000) discusses this issue.

sent account is the manner in which it explains the representational intentionality or aboutness of perception of an object X by an organism Z. It involves the relevant *reflexive causal interaction* between Z and X, in which X causes Z to modify, not just any of its causal dispositions, but specifically those toward X itself — *a reflexive causality*, or simply a *reflexive* theory of perception.

More on the Reflexive Theory of Perception

The current reflexive theory of perception will now be explained in more detail.[5] Recall that the basic idea is that a genuine perception of object X by organism Z involves both the relevant perceptual state S of Z being caused by X, and that the state S involves *changes in Z's X-related dispositions*, activation of which dispositions would provide causal evidence that Z has perceived X, over and above any initial evidence that state S was caused by X.

As for the first condition, as remarked above it is generally agreed that one necessary condition of an organism Z perceiving object X is that object X *causes* the relevant perceptual state of X, in the minimal sense that it initiates the causal chain that results in the perception of X. But there is an important intermediate step in that causal chain, namely a *peripheral sensory state z1* of Z, which state constitutes the immediate physical reception of X's causal effects on Z, such as the state of retinal stimulation caused in Z's eyes by light from X. It is peripheral states such as z1 that were argued not to be carriers of worldly information. Thus, to summarize so far, one necessary condition for Z to perceive X is that X causes such a sensory state z1 in Z.

On the present account, such peripheral sensory states z1 in turn cause further internal perceptual processing to occur, resulting in some distinct internal state IS, that is the actual physical basis of Z's perceptual representation of X. Hence the second necessary condition for Z to perceive X, or for state IS to perceptually represent X, may be stated as follows: the requirement is that sensory state z1 causes Z to acquire some X-related disposition, where Z's state of acquiring the disposition may be identified with the relevant internal perceptual state IS. Thus internal state IS is such that its causing by object X — via X causing state z1 to cause IS — results in *changed* or modified *dispositions* for Z to causally relate to that object X, relative to the dispositions it would have if it were not perceiving X. This condition could be called the *reflexive disposition change* condition on perception. It will now be given a preliminary analysis.

First, the condition is *dispositional*, so as to avoid a crude behaviorism that would require perception of X to always involve immediate or concurrent

[5]For simplicity, only correct cases of perception will be considered in this initial account; misrepresentation issues are briefly discussed later.

causal interaction with object X, or avoidance of such. But the concept is intended to be inclusive enough to include activated dispositions — actual behavior caused by the disposition — as well as inactive dispositions. Second, it is left open whether the relevant dispositions are broadly *functional* ones, in the sense that they could be identified independently of any particular causal means by which a concrete kind of interaction with X, such as moving X or eating X, might be carried out during any actualization of the relevant disposition, or whether they are identifiable only purely causally.

Third, the *change* or modification aspect is needed so as to distinguish any previous non-perceptual causal relatedness to object X by Z, such as Z's having physically grasped X without having had any perceptual awareness of having done so — even if the proximity of X ensures that internal state IS is caused by X — from the changed X-related abilities available when Z does perceive that it has already grasped X. And fourth, it is unnecessary to require an additional condition of a conceptual or cognitive kind on perception of X, such as a claim that state IS must involve a *classification* or *categorization* of the relevant cause of the perception as being an object of kind X. For on the present account, the cognitive implementation of such higher level cognitive processes would itself presuppose the availability of lower level perceptions of an item X as being an X. Thus, for example, there could not be a general-purpose sensory classification module pre-wired to recognize Xs, since perceptual recognition of Xs requires the right kinds of connections to exist between any internal module and the organism's X-related activities, absent which the output of the module would not count as a perceptual recognition of an X, and hence not as a correct classification of it as an X either.

To summarize, there are two necessary conditions of perception of X, or for state IS to perceptually represent X, according to the current reflexive view of perception. The first is a *causal* condition, that a peripheral sensory state z1 of Z must be caused by X, and the second is a *reflexive disposition change* condition, requiring that state z1 cause a perceptual state IS of Z, which state is the acquisition of some X-related disposition. Also, these two conditions are jointly sufficient for perception or perceptual representation of X as well.

But what of perception not just of objects X, but properties F of X also? At the purely causal level there is no principled distinction between objects and properties, so exactly the same analysis applies. For person Z to genuinely perceive the color F of an object X, that color — or the purely physical properties that constitute its physical realization — must themselves cause sensory mechanism z1 of Z to cause Z to acquire some color-related disposition with respect to X, such as one manifested by Z sorting X into an appropriate box based on X's color.

As for the issue of perceptual representation itself, this is as close as one could hope to a purely naturalistic, low level and minimalist analysis of the

concept. The reflexive view imposes or presupposes no further requirements of similarity, intentionality, having a content, informativeness, and so on. Later discussion will show how the reflexive view may be extended to include conscious perception and perceptual content.

To summarize the last two sections, the reflexive view of perception potentially provides a completely naturalistic view of perceptual learning about the world, as involving changes in the causal dispositions of an organism Z with respect to the objects X that it has previously perceived, and more generally with respect to any perceptually similar objects. Since there are no inherent limits on the causal powers and complexities of these changed dispositions toward worldly objects, and since evolutionary considerations can explain why current species are generally successful in their perceptual learning, a reflexive view of perception is potentially able to explain cognitive concepts, such as those of acquisition of justified belief and knowledge, without any need to appeal to naturalistically suspect and causally idle informational concepts.

Thus the relevant reflexive causality concept is not itself an informational concept in the usual senses, and so it is possible to maintain that, strictly speaking, there is no information of any kind in the universe, whether worldly information, or information internal to organisms capable of perception. Hence whatever truth there may seem to be in informational views should be explained away as claims about the X-caused, modified X-related dispositions, and hence of the resultant abilities or skills, of the relevant organisms Z. The concept of information would then persist only as an informal or folk-psychological preliminary to such dispositional claims.

Worldly Information as Bad Science

Some additional arguments, going beyond those provided above, will now be presented against the introduction of a concept of information for purposes of scientific explanation. To begin, here is an evolutionary argument against the idea of perception as always involving acquisition of information. Even if an organism *could* "acquire" typical items of information, such as that object X has some property F, such acquisitions have no clear connection with evolutionary advantage, since the relevant information might be completely useless to, or even inimical to the survival of, the relevant species engaging in such informational practices. The idea that the selective survival advantages of perception to a species may be explained in terms of information intake is specious, since both winning and losing species would equally be taking in information of some kind. By contrast, the reflexive theory makes no corresponding assumption that perceptually changed dispositions necessarily have a positive effect upon an organism's learning. All that could possibly matter, from an evolutionary perspective, is that species and their members

should be able to use perception to modify their behavior in adaptive ways—an issue directly addressed by the current reflexive view, but ignored by informational views, since changes in quantity or kind of information internally stored by organisms have no clear relevance to survival.

A related epistemic argument against the worldly information view is as follows. Whether or not theorists demand purely naturalistic standards of explanation, the postulation of worldly information is empty or idle in any case, because parallel, purely causal explanations are available, which perform the same epistemic explanatory work without any need to postulate extra informational entities or properties.[6]

Recall that assumptions of worldly information have probably been based on some rough intuitive argument of the following kind. Perception of object X by organism Z is normally the result of object X causing the relevant peripheral perceptual state S in Z, such as a concrete state of retinal stimulation. If this process results in Z *acquiring new perceptual information* X' about X, then that new information X' must have been acquired by Z via the relevant causal process. But since the information acquisition happened purely as a result of the relevant retinal stimulation S, worldly state S must itself have provided the new information X' about X. The inadequacy of this argument should be clear from the discussion in the first section. But now the current epistemic point becomes relevant, namely that the postulated worldly information transfer, via the retinal stimulation, is *not needed* to explain the relevant perception and perceptual representation.

Here is roughly how a reformulated argument would go, using the two necessary conditions for the reflexive view of perception. To acquire a changed ability to interact with object X, organism Z must have internal structures that reliably connect peripheral sensory states z1 of Z to X-related dispositions. The evolution and maintenance of such structures would not have been, and would not be, possible without equally reliable causal correlations between an object X and those external states z1. For example, in evolutionary terms, the connections between X and z1, and z1 to the relevant X-related dispositions, must be stable enough so that evolution via natural selection can take place.

Thus standard reliabilist kinds of epistemic argument may be used to explain perceptual successes.[7] So, with respect to the initial, causally reliable connections between X and z1, they provide background necessary conditions of perception of X, as expressed in the initial cause-by-X necessary condition in the present reflexive analysis, but any additional claim that they

[6]A point which Dretske (2000) now acknowledges for kinds of perception not involving learning, though without any more general concession.

[7]E.g., Goldman (1976), Kornblith (2002).

provide necessary *information* to Z, that somehow is supposed to explain its modified, perception-based abilities, is completely gratuitous. Indeed, arguably such views simply miss a major point of evolutionary explanations, which postulate built-in genetic competences to organisms so that they can directly exploit causal regularities, rather than having to learn to do so via isolated items of information, one at a time.

Statistical Information Transmission as Bad Science

Millikan (2004) argues that there are "local natural signs," such as a footprint, that is a sign of the local presence of some animal X. She argues that, though in her view Dretske's strong concept of information — involving a probability of 1 — is unworkable, nevertheless there remains an important concept of *statistical information*, such that the footprint provides some positive statistical information to perceivers that there is a significant probability that an X is in the offing. Indeed, in her view the perception of such a sign can itself amount to a kind of perception of the animal X of which it is a sign, since on her analysis all perception is mediated by signs or representations, whether internal or external.

However, such an exotic example of statistical information will be replaced by a simpler, structurally similar one for present purposes. On Millikan's account, presumably seeing an X in the normal way (not via external signs), but under adverse perceptual conditions, would involve a similar statistical transfer of information to a perceiver, with a probability of less than one. But a significant problem with such an account is that perception seems to be an all-or-nothing matter: one cannot partly, or with some statistical probability, perceive a particular object X on a given occasion. Thus a more plausible analysis of the statistical claim is as a long-run frequency claim, such as that on 100 occasions of attempted perception, X was perceived 70% of the time.

But then the frequency does nothing to explain how any *particular* such perceptual episode does, or does not, involve a perception of the X. Thus a claim that any particular such episode of successful perception of X could be explained by the transmission of the statistical information to the perceiver is inadequate, since presumably the unsuccessful cases would involve transmission of exactly the same statistical information to the perceiver, which information therefore cannot differentially explain why some of the cases were successful but others unsuccessful.

As for the present reflexive account, it has no problem handling such statistical cases. To begin, it simply refuses to accept the supposed epistemic authority of the statistics in explaining particular cases. It can insist that the relevant statistical probability is epistemically no more than a generalization from the percentage of successful cases of perception of X in a given situation.

Why some were successful and others not is a matter for further investigation of the *whole* relevant causal/dispositional chain from X to Z's enhanced abilities (or lack thereof) in each particular case; but it is gratuitous to claim that, for instance, the failed cases must have occurred because of the "informational weakness" of the causal connection between X and peripheral state PS, or the successful cases because of its statistical strength.

To be sure, presumably there are characteristic differences in resulting peripheral states PS in cases of attempted perception of X in ideal versus adverse circumstances, but presumably there will also be characteristic differences in the corresponding internal states IS, and resulting characteristic differences in acquired X-related dispositions as well. Thus statistical informational accounts of perception under adverse conditions are simply bad science, which do not consider *all* of the relevant data in perceptual failure cases, and which reify probabilities as if they had some epistemic worth independent of the actual causal factors in particular successful versus failed cases of perception.

On the other hand, insofar as some genuine learning could take place in such cases, the reflexive view can readily explain it. If genuine perception of X by organism Z occurs in 70% of the attempted cases, then in those positive cases Z's X-related dispositions will be appropriately changed, even if there are no such changes in the other 30% of failed cases. Thus if Z is hungry for Xs, and by attempting to perceive them is successful 70% of the time in eating them as a result, then Z will have been well fed via the successful, changed-disposition cases, in spite of receiving no perceptual information whatsoever about those Xs in any of the cases, whether of a statistical kind or not.

Misrepresentation Issues

As a preliminary point, one should not overstate the significance of issues of perceptual misrepresentation. A world in which no misperception whatsoever occurs is at least empirically possible, even if our world is not such a world, whereas semantically speaking, presumably both true and false propositions must be available in any possible world. Thus it is a mistake to assume that a theory of perception *must* assign significance to cases of misperception equal to that of correct perceptions, even if any adequate semantic account of propositions must assign an equal role to both true and false propositions. An idealized theory of perception that ignored misperception issues could be a legitimate scientific theory, just as much as a dynamical theory that ignores issues of friction. Nevertheless, it is a standard complaint against informational views of perception that information as such is necessarily correct or true, so that such views could give no adequate explanation of the kinds of perceptual misrepresentation that actually occur in our world, nor, more generally, of the fundamental bipolarity of semantic and propositional concepts. However,

the current reflexive view has powerful resources for explaining typical misperception cases that are unavailable to the information theorist.[8]

Hence, though a reflexive theory of perception can easily explain cases of misperception, the importance of this success should not be over-estimated. Indeed, as a non-semantic, non-cognitive and non-conceptual account of perception, the reflexive view can claim to avoid in any case the main thrust of misrepresentation-based criticisms of informational views, which clearly are intended to provide foundations for bipolar semantic and cognitive concepts (Cummins, 1996). [At this stage it is an open question as to how a reflexive view of perception would relate to such higher level concepts.] But, with those reservations noted, it is clearly a significant advantage of an empirical perceptual theory if it can explain real-world kinds of perceptual mistakes that typically occur, and do so in the same theoretical terms in which it explains cases of correct perception, rather than as mere approximations to, or degenerate cases of, its main theoretical constructs.[9] This the reflexive view can do, as follows.

Recall that the reflexive view specifies two necessary conditions for perception, or perceptual representation, of an object X to occur. There is a *causal* condition, that X must cause the relevant peripheral sensory state z1, and also a *reflexive disposition change* condition, that state z1 in turn must cause some change in Z's X-related dispositions, relative to those that it had prior to the start of the current perceptual episode. Now as an initial point showing the compatibility of this analysis with misperception, there is no built-in requirement of the reflexive view that the exercise of abilities or dispositions by organism Z would necessarily result in *correct* X-related behavior. Further, even if a desirable scientific concept of disposition-based abilities did presuppose generally correct exercises of them, the reflexive theory only requires a change in dispositions for perception to occur — which change might result in the loss of some previous abilities in spite of a gain in others.[10] So in sum, the *reflexive disposition change* condition of a reflexive theory imposes no inherent conceptual restrictions on the possibility of misperception, unlike the basic, necessarily correct theoretical construct of an information-based theory of perception.

For example, there are several ways in which organism Z might be perceptually caused by an object X to change its abilities with respect to X's *shape*,

---

[8]For a complementary account consistent with the reflexive view see Dilworth (2005a, 2005b).

[9]Also, over a longer time line, in terms of evolutionary theory, the existence of less than optimal cases of perception is presupposed in the assumption that some species have perceptual skills inferior in practice to those of other species, that result in selective survival disadvantages for them.

[10]See Millikan (2000) for a useful discussion of such distinctions between abilities and dispositions.

size or *position*, some of which we would describe as being correct and others incorrect. If the relevant ability or disposition changes in Z with respect to shape, size or position were actualized as attempts by Z to manipulate X by moving or grasping it, then in general correct perception would be shown by success in the manipulation attempts, whereas failure, such as shown by grasping movements in the wrong place, or with insufficient closing of a hand or arms, would provide evidence of misperception with respect to the relevant properties of X — which evidence, it will be recalled, is behavioral evidence of the existence of a changed, X-related disposition or ability caused by X.

Intentionality and Perceptual Content

It is at least convenient, if not theoretically indispensable, to discuss higher forms of perception — as well as higher forms of cognition generally — using a broadly intentional vocabulary, including such concepts as those of intentionality, aboutness and content. Here some ways are suggested in which such discourse could be shown to be consistent with the current, strongly naturalistic reflexive perceptual theory.

Recall that specifically information-based concepts of representation, aboutness or content — as based on causal covariation or statistical correlation of some kind — have no legitimate place in a scientific perceptual theory. The suggestion is that no other scientifically or metaphysically realist theories of those concepts are viable either, so that the options for retaining perceptually relevant intentional concepts must instead be found or produced via reductionist, eliminativist or irrealist theories of those concepts. In this section it is shown that specifically perceptual aboutness and content are physically reducible, while in the next section broader forms of intentionality involving intentional objects will briefly be given an irrealist construal.

To begin, the basic reflexive view arguably already involves a naturalistic reduction of the intentional concept of *aboutness*, in the case of actual objects of perception. The relevant X-caused abilities or dispositions in an organism Z are "about" X in the sense that they are typically activated as immediate behavior by Z that is X-related, or about X, in a purely causal sense, involving actual causal interactions of Z with object X, or avoidance thereof. So the relevant concept of representational X-aboutness is, at the lowest level, typically concretely exemplified by organism Z in such actual X-related activity.

As for higher, conscious levels of perception, any of the dispositional changes caused by X could be typically manifested as immediate X-related activity of some relevant kind, so that the consciously episodic, rather than dispositional, aspects of perceptual events could be explained in such occurrent activity terms. Also, arguably the specifically conscious aspects of per-

ception are to be explained in terms of broad monitoring and control activities with respect to the relevant dispositional changes, so that the conscious aspects of perception summarize or model the relevant factors involved in the relevant changed, X-related dispositions, so as to achieve improved control over them.[11] For example, the dispositional changes involved in a perception of the shape of object X would presumably be consciously monitored by the cognitive construction of a depictive model of that shape, any inconsistencies or other problems with which would prompt distinct but related perceptual contact with object X, such as by viewing it from a different angle.

In the higher, conscious levels of perception, typical relevant X-related dispositions would presumably involve immediately activated dispositions to *model* X and its properties, with the model having a structure similar to that of X itself, and with each part of the model involving an appropriate group of X-related dispositions. For example, a modeling of the various differently shaped areas of a perceived object would involve similarly structured groups of dispositions to causally interact with each of those areas in distinctive ways, that depend on what each area of the shape is modeled as being. Nevertheless, such a model would still be *about* X, not because the model has the same structure as X (since many other objects would also have the same structure), but because each part of the model structure involves appropriate X-related dispositions. Thus the relevant activated modeling dispositions, as well as the resultant X-related dispositions associated with each part of the thus-constructed model, would be an integral, controlling and monitoring part of the whole array of dispositions that constitutes organism Z's *complete* current perceptual representation of object X.

Now the relevant model models how object X *seems* to be to organism Z, rather than how it actually is, because some of the modeled properties of X may not be its actual properties, but instead misrepresentations of those properties. Thus, as one would expect, the relevant model would misrepresent X in ways closely related to those more directly causal ways in which Z misrepresents X. For example, a misrepresented shape in the model would involve dispositions to grasp or avoid X in ways that would also provide evidence of incorrect shape-related perceptual processing by Z.

As for the *content* of the relevant perceptual state, it may be identified with the properties that the relevant model *models* X as having. This is a more specific view than the common view of perceptual content as the properties that an object X is *represented* as having (e.g., Tye, 2002), but it is closer to the usual understanding of content properties as those that are consciously, or at least cognitively, accessible to the perceiver, whereas the global

[11]For a recent account of the broad scope of such control activities see Clark (2001).

reflexive view of "properties an object is represented as having" includes also low level dispositions that presumably are not even cognitively accessible.

But in any case, whatever the appropriate range of content properties may be, that issue is a relatively peripheral one for the reflexive theory, since it explains *all* specifically perceptual ways of representing an actual object X in purely physicalistic, causal and dispositional terms. Thus, to state the matter explicitly, a reflexive theory claims to achieve a naturalistic reduction of perceptual content properties to purely physical properties of the perceiver, in all cases of perception of actual objects (omitting hallucinations, after-images and so on).

To be sure, this view does not immediately imply that *all* content properties, including cognitive, linguistic and consciousness-related properties generally, are purely physical, since it might be arguable that the modularity of perception (Fodor, 1983) permits physicalistic explanation in its case, but not necessarily more generally in explaining the whole range of cognitive architecture. But at the same time, if aboutness and intentionality are thus reducible for all perceptual properties, including conscious ones, broader reductions may well be feasible. The potential virtues of an eliminative irrealism for some intentional objects will be presented below, hence potentially facilitating such a broader reduction.

Intentional Objects and Perception of Pictures

There is at least one class of perceivable objects that are potentially threatening to the current, purely physicalistic reflexive account of perceptual content, namely representational pictures such as photographs or paintings. One could see a concrete painting X of a man Y, who is represented by the physical picture, but in seeing that man Y in the picture, one's perceptions are not directly caused by an actual man, even if the picture does represent some actual man — because, of course, the picture is not itself an actual man.

This is threatening to the present account in two ways. First, it seems to be a case of perception of an object Y in which there is no directly perceptually available object Y to cause the perception. And second, since object Y is not directly perceptually available, it is unclear in what sense one can be "representing Y as having certain properties" in the usual dispositional sense provided by the reflexive theory, in that these would seem to be, at best, dispositions directed toward a perceptually absent object.

The situation is made worse by the possibility that there is no such man Y represented by the picture at all, since it is commonplace for artists to paint generic subject matters, such as lakes or people, without there being any corresponding actual subjects that the pictures represent. In such cases, there

cannot be any actual man Y toward which the perceiver's representational dispositions are directed as part of his perceiving the man Y in the picture.

The outlines of a solution are as follows.[12] The case may be analyzed as one of *perceptual dependency*, in which some of the X-related dispositions generated during perception of the actual, physical painting X, which initially are purely about X itself, may be *repurposed* or redirected toward a purely virtual object Y. Thus in perceiving a picture, such as a painting X of a man Y, one "perceives" the man Y in an extended sense, in that the physical picture X is so constructed that the perceptual states it causes in viewers involve dispositional changes that strictly result in misperceptions of X itself — since the picture X is not itself a man — yet which dispositions have a coherence and integrity *similar* to those of dispositions involved in veridical perceptions of an actual man. Thus in such a case, in virtue of those similarities a perceiver of the picture initially acquires dispositions toward a *virtual* man Y — since what is seen on the canvas is not of course itself identical with an actual man.

Now suppose that the picture X is a picture *of* some actual man Y, such as Einstein. Then its representing Einstein can be explained in terms of those same dispositions toward a virtual object Y, among perceivers of the picture, amounting in this case to *indirect actual* dispositions toward Einstein himself (such as to read more about him, or further his theories). On such a view, pictures and other public representations play a role as intermediary or stand-in objects X in the service of a kind of *indirect* perception, by viewers of X, of the actual objects Y that they represent, though of course these are not genuine or direct cases of perception of the object Y.[13] It is a feature of this account of pictorial representation that, whether or not there is some actual entity Y represented by picture X, there will still be the same subset of basic dispositions, causally created in viewers of that picture X by their viewing of X, toward a *virtual* object Y — i.e., even if those concrete dispositions can also be regarded as indirect dispositions toward an actual object Y represented by picture X.

With respect to that virtual object Y, it may be given an *eliminative, irrealist* construal — the relevant subset of repurposed, Y-related dispositions may be construed as being "about object Y," even though in actuality there is no such object. Thus the explanation given is, in more general terms, that the specific intentionality or aboutness of perceptual dispositions depends on how they are characteristically used — and in the case of pictures, their characteristic uses include such virtual-object aboutness cases with respect to their depicted subject matters. Thus, in sum, if perceptual content is taken to include both represented objects and the properties that they are represented as having,

[12]See Dilworth (2004) for a more comprehensive account.

[13]See Dilworth (2004, Ch. 11).

the overall reflexive proposal is that represented properties as such are always physicalistically reducible to causal dispositions, while represented objects may in "virtual" cases be analyzed away via an eliminative, irrealist construal.

A similar approach would seem at least initially promising in other more paradigmatic intentionality cases involving possibly non-existent objects — hence also showing that the above irrealist construal is not merely an ad hoc response to purely perceptual cases. For example, the fact that someone could search for, or hope to find, explorer Ponce de Leon's legendary "fountain of youth" could appropriately be explained as a case of actual searching-related dispositions toward an object that may or may not exist in actuality. In either case, the relevant dispositions are in the first place toward a purely virtual object, which dispositions may or may not (depending on actual facts about the universe) turn out to be also indirect dispositions toward some actual object. Thus similar concepts of representation and content could be applied to such non-perceptual cases.

On Directness Versus Indirectness of Perception

One interesting remaining issue about the reflexive view of perception is as follows. Consider the traditional issue as to whether perception puts us into direct contact with worldly objects and properties, or whether instead some kind of internal perceptual intermediaries are involved in perception. Because of the dispositional, purely physicalistic nature of the reflexive theory, neither traditional alternative makes complete sense as applied to it. As an initial point, the theory might seem clearly to be a direct theory of perception, in that a perception of object X is directly caused by X, and it is X itself, the actual physical object, to which representational properties are attributed in perception — via the X-related dispositions. However, at the same time, a disposition as such involves no worldly contact at all, so in that respect it seems more like an indirect perceptual intermediary, since it is only *activation* of a given disposition that directly affects the actual object X itself. But since not all of the changed dispositions making up a perception of X need be activated immediately, some components of a given perception may be indirect inactive dispositions, while others are direct activated dispositions.

As for the conscious, activated and occurrent modeling of X that in part presumably makes up the conscious part of perception, it too has an equivocal indirect/direct status. Insofar as it involves an activated disposition to model X, its results might seem similar to a traditional indirect element in perception — a kind of pictorial intermediary. However, on the reflexive view that modeling only occurs in order to set up additional, structured X-related dispositions that also have a role in controlling the lower level, more directly behavioral X-related dispositions, so in that respect the constructed

model seems more directly connected to object X itself. Thus, to conclude, the reflexive theory potentially provides an independent, purely naturalistic perceptual theory that deserves to be evaluated on its own terms, whether with respect to the traditional direct versus indirect distinction, or more generally with respect to its role in replacing informational approaches to perception.

References

Agassi, J. (1966). Sensationalism. *Mind, 75*, 1–24.
Clark, A. (2001). Visual experience and motor action: Are the bonds too tight? *Philosophical Review, 110*, 495–520.
Cummins, R. (1989). *Meaning and mental representation*. Cambridge, Massachusetts: MIT Press.
Cummins, R. (1996). *Representations, targets and attitudes*. Cambridge, Massachusetts: MIT Press.
Cummins, R., and Poirier, P. (2004). Representation and indication. In P. Staines and P. Slezak (Eds.), *Representation in mind* (pp. 21–40). Dordrecht, Holland: Elsevier.
Dilworth, J. (2004). *The double content of art*. New York: Prometheus Books.
Dilworth, J. (2005a, in press). The double content of perception. *Synthese*.
Dilworth, J. (2005b, in press). The twofold orientational structure of perception. *Philosophical Psychology*.
Dretske, F. (1981). *Knowledge and the flow of information*. Cambridge, Massachusetts: MIT Press.
Dretske, F. (2000). *Perception, knowledge and belief*. Cambridge: Cambridge University Press.
Fodor, J. (1983). *The modularity of mind*. Cambridge, Massachusetts: MIT Press.
Fodor, J. (1990). *A theory of content and other essays*. Cambridge, Massachusetts: MIT Press.
Goldman, A. (1976). Discrimination and perceptual knowledge. *The Journal of Philosophy, 73*, 771–791.
Kornblith, H. (2002). *Knowledge and its place in nature*. Oxford: Clarendon Press.
Lowe, E.J. (2000). *An introduction to the philosophy of mind*. Cambridge: Cambridge University Press.
Millikan, R. (2000). *On clear and confused ideas*. Cambridge: Cambridge University Press.
Millikan, R. (2001). What has natural information to do with intentional representation? *Philosophy, 49* [Supplement], 105–125.
Millikan, R. (2004). *Varieties of meaning* [The 2002 Nicod Lectures]. Cambridge, Massachusetts: MIT Press.
Tye, M. (2002). *Consciousness, color and content*. Cambridge, Massachusetts: MIT Press.