PHILOSOPHICAL TOPICS
VOL. 33, NO. 1, SPRING 2005

Disjunctivism and Perceptual Psychology

Tyler Burge
University of California, Los Angeles

Perceptual anti-individualism holds that the nature and correct individuation of perceptual states and perceptual beliefs are constitutively associated with relations, including causal relations, between capacities in the perceptual system and aspects of the physical environment.

A closely associated thesis is that

> a constitutively necessary condition on perceptual representation by an individual is that any such representation be associated with a background of some *veridical* perceptual representation.

For purposes of this essay, in accord with perceptual anti-individualism, I will assume that the relevant veridical perceptual representation is representation of entities in the physical environment.

Perceptual anti-individualism—sometimes called *externalism* about perception—is very old, stemming as it does from Aristotle. In the history of philosophy it has been the dominant view. In the last quarter-century it has become widely held in an explicit way. I believe that it is presupposed and relied upon in perceptual psychology. I think that it and its associated thesis are surely true.

Proponents of perceptual anti-individualism differ over the natures and correct individuation of perceptual states and perceptual beliefs. Some differences can be settled by reflecting on other things that we know. Intra-psychological relations, as well as relations to the environment, constrain the natures of representational

explanation, and cognitive ability conspire to force recognition of the two kinds of representation mentioned in the preceding paragraph. In the Appendix, I criticize attempts to support disjunctivism.

I. PERCEPTUAL ANTI-INDIVIDUALISM

The key feature of perception is that it is representational. I mean this in a very broad sense. Perception is representational in that its nature is both to purport to be about something and to represent it as being a certain way. To purport to be about something is to *function* to represent something—to function to refer to or indicate something.

Perception's function to represent something is fallibly realized. When the function is successfully realized, perception represents particulars and represents them veridically as being a certain way. Being fallible, perceptual representation may not succeed in representing a particular as it is. It may even fail to represent any particular at all. The exercise of perceptual capacities can be veridical or non-veridical—accurate or inaccurate. Perceptual states are typed in terms of their representational content—in terms of how they represent things to be.

Representational content is an abstraction that marks the kind of psychological state that has the content, and that is what is true or false, veridical or unveridical. For example, the representational content of a belief that most dogs are faithful is the abstract thought content that most dogs are faithful. I will discuss the representational content of perceptual states shortly. Roughly it is a percept-abstraction that attributes or categorizes what is purportedly perceived as being a certain way (as being round, for example) and that purports to single out a particular instance of that property or relation, or purports to single out a particular that has that property, or enters into that relation. All elements of representational content function to represent—all are representational.

I believe that there is a significant distinction between the forms of representational content in perception and those of thought. These forms *mark*, or help type-identify, psychological abilities whose operations are structured in different ways. Thought is propositional. Concepts are certain representational constituents of propositional thought contents. Perception is representational but non-propositional. Primitive perception has no conceptual elements. The difference between perception and thought lies in the type of representation, the way representation is organized, and the abilities marked by the representational content. With some serious qualifications, I believe that perceptual representational content has a structure more like a map. Thought has propositional form, of the sort exemplified by sentences. What I say here will not depend on taking perceptual representation to be non-propositional and non-conceptual. Anti-individualism applies to both perception and empirical thought about the physical environment.[2]

The full representational content of a perceptual state is normally very complex. It is common to discuss parts of this full content. For vision, a representational content of a perception of a scene includes representational contents that specify parts of the scene, or properties of objects in the scene, or objects in certain relations in the scene. We can think of a perceptual state as the state of representing the whole scene. Or we can think of a perceptual state as representing some sub-part of it. Either way, the perceptual state is type-identified in terms of its representational content.

Perceptual representational content types *mark* or help type-identify perceptual states and abilities. These states and abilities involve perceptual response to types of kinds, properties, and relations that figure in the individual's needs and activities in its normal environment. Perceptually responding to particulars with certain shapes and colors that are relevant to the activities of eating, avoiding being eaten, reproducing, navigating, finding shelter, and so on, is a fundamental function of the visual system. The representational content of an animal's perceptual states is individuated partly in terms of what causes those states and how those states enable the animal to cope with specific types of entities in its environment. Successful interactions help ground individuation of perceptual states partly in terms of representational contents. These contents in turn mark or help type-identify representational abilities. The psychological kinds or natures thus individuated enter into a pattern of animal activity that is the subject of psychological explanation.

This is the basic shape of perceptual anti-individualism. The shape is, however, more complex than it may initially appear. I mention two sources of complexity.

An individual can be systematically mistaken in application of its perceptual representations if it is in circumstances other than those that the perceptual system functions to represent. If a frog or a child is given a show of holograms, its perception may fall into systematic error. If it is moved into a hall of prisms, mirrors, and special lighting, it may be radically fooled. These points reflect the fact that perceptual contents and abilities are often typed not in terms of the individual's history, but in terms of successes that serve the basic needs of perceivers with the relevant perceptual system.

The point about error goes further. An individual can be perceptually wrong more often than right even in its normal environment. The value of veridicality may pay for many errors. One veridical representation of a predator may pay for a lifetime of false positives. The key idea is that the individuation of perceptual types is ultimately *explained* in terms of conditions for veridical cases.

A second source of complexity is that perceptual anti-individualism does not require that every perceptual state type sometime be caused by instances of what it represents. Perceptual representation can involve spectra containing many contents that are never satisfied by particulars perceived and never caused by instances of the property represented. Consider a perceptual system capable of representing shapes. Many of the more complex shapes may never have been instantiated in the environment that the perceptual system functions to deal with. A given shape representation might be triggered only in cases of illusion—even in the individual's normal

otherwise derived from any actual scene. The photograph represents an actual scene containing objects and their properties. The painting represents types, but does not even purport to represent particular, actual instances of those types.

Singular representational elements in perception are often neglected.[6] I shall return in section VII to a more detailed discussion of them.

The second aspect of the representational content of perceptual states concerns its *perspectival nature*. Representational perceptual content is to be strictly distinguished from the entities (objects, properties, or relations) that are perceived. Such content always constitutes a partial representation of the particulars perceived or the properties or relations attributed. For any given particular object, property, or relation, there are many possible (commonly actual) representational contents that correctly represent it. Both singular and general elements are perspectival.

An aspect of the perspectival nature of representational perceptual contents that is salient in human perception is that phenomenologically different contents can apply (and can be taken to apply) to the same particular, property, or relation. The general phenomenon of perspectival representational content is ubiquitous in perception—phenomenally conscious or not.

A perceptual state's representational content is fundamental to the kind or type of state that it is. The content that marks or helps type-identify perceptual states is, in the first instance, not the referent, but the mode of presentation.[7] It depends on the individual's perspective and marks abilities that are exercised from that perspective. They are abilities to perceive that connect to a perceptual referent only from some perceptual perspective, only in some context. Perceptual representational content marks such perspective and context.

Suppose that the perceptual system or the animal's behavior allows for the *possibility* that two occurrent perceptions at different moments are referentially associated with different represented entities. These may be properties or relations, types or tokens, particulars or kinds. The individual and perceptual system may be perceptually referring to only one entity. They may even treat the different perceptual states *as* referring to one. Still, if there is a psychological/logical possibility that the perceptual states are perceptions of distinct entities, the representational content—the mode of presentation—varies. The individual's perception can correctly treat a perceptual referent as the same, even though the identification is not guaranteed by the form of the perceptual content, or the perceptual abilities marked by the content. Then we have an analog of Frege's true identities formed by terms expressing different modes of presentation.

For example, suppose that the system treats the referent of a touch as the same as that of a visual perception. Or suppose that the individual perceptually tracks an object as it turns. The perceptual representations will commonly have different representational content, even if the contents in fact do have the same referent, and even if there is a further intermodal, or modality neutral, representational content. There is the possibility of error in the identification. Non-trivial psychological mechanisms enter into the identification.

Of course, when one fails to perceive something, there is perceptual representational content but no perceptual referent. For these reasons and others, perceptual representational contents, hence perceptual states, are not individuated purely in terms of the environmental referent.

III. THE EMPIRICAL PSYCHOLOGY OF VISION

I believe that there are general arguments for perceptual anti-individualism that do not rely on any specialized knowledge. The arguments make use only of general, uncontroversial, well-known empirical facts, and other rational considerations. I believe that perceptual anti-individualism provides the only acceptable framework for understanding conditions under which perceptual representation is possible. And I believe that one can recognize this on reflection and by considering alternatives. I will not discuss these general arguments here. Here I discuss a body of considerations centered in concrete empirical explanation.

Perceptual anti-individualism is embedded in the practice of the empirical psychology of perception. Empirical psychology takes for granted the general anti-individualist account. It provides empirical explanation of the specific, contingent ways that perceptual-state kinds depend on relations to the physical environment. And it makes uses of perceptual-state kinds which are embedded in law-like generalizations that depend on such relations.

Healthy modes of interchange between philosophy and science cannot be reduced to a formula. Philosophy's record regarding perceptual psychology in the last century has not, however, been exemplary. A good deal of philosophy has proceeded with insufficient reflection on the science, or has offered unconvincing rationales for taking it to be irrelevant to philosophical problems.

The psychology of perception, particularly vision, has become serious science. It has well-established results and successful application of mathematical methods. There is no good reason to doubt that it provides insight not only into the mechanics of perception, but into aspects of its nature.

I will go over some basics of the psychology of vision. These will illustrate how perceptual psychology embeds and gives empirical specificity to the general anti-individualist account.

The theory of vision begins with the observation that detectors in the retina are sensitive to the effects of arrays of light frequencies. Its paradigmatic problem is to explain how perceptions of the distal environment are formed from sensitivity to such light arrays—registration of their spatial and temporal distributions on the retinal detectors. There are other sources of input into the visual system—proprioceptive input, input from other senses, top-down cognitive input. Still, as an empirical matter, it has been repeatedly confirmed that many basic explanations of fundamental visual processes can be successfully carried through while bracketing these further sources, factoring them in at further stages of explanation.

Underdetermination has been shown empirically to take an immense variety of forms. Non-technical, intuitive considerations, however, illustrate the basic fact of underdetermination.

Ambiguous figures certainly suggest underdetermination. The Necker cube or the duck-rabbit drawing bring out the role of the visual system in producing a state that is "committal" beyond what is present in the proximal stimulation itself. Of course, in the standard drawings, the external two-dimensional object is the same, however it is perceived. Nonetheless, the drawings suggest the possibility of different objects producing the same visual stimulation by showing that the same visual stimulation is compatible with different perceptual representations.

Visual illusions illustrate the point more directly. The Ames room is a trapezoidal room with a sharply receding back wall. From a certain perspective, it is mis-perceived as being rectangular; and the sizes of familiar objects (human bodies) in it are also mis-perceived because distance relations are mis-perceived. The same visual stimulation could have been produced by a scene that made the same perceptual-representation types veridical. The proximal stimulation is compatible with either of these two possible objective situations. This point applies to numerous brute perceptual illusions.[11]

A further intuitive consideration that illustrates underdetermination is what is known as visual completion. In cases in which one object occludes another, the occluded object is perceived as continuous. In cases in which the front of an object occludes its back, the object is perceived as a body, and often it is perceived as having a particular three-dimensional shape. In the first case, the proximal stimulation is often compatible with there being no occluded object—only two objects adjacent to a middle object. In the second case, the proximal stimulation is compatible with the object's being a mere facade, or having any number of oddly shaped backsides.

Perhaps the most basic intuitive consideration illustrating underdetermination lies in reflection on the geometrical consideration noted earlier. The light intensities that constitute the proximal stimulation are registered on the retina in a two-dimensional array. The array corresponds to a physical array in the receptors—each corresponding to a surface area of stimulation. The information that is registered can be constructed as a two-dimensional image. There is a determinate solution to how light from a three-dimensional scene will project onto a two-dimensional surface. The visual system must, however, use the two-dimensional array of information —the light intensity that stimulates each unit surface area of the retina—to perceptually represent a three-dimensional scene. This "inverse problem" has an infinity of mathematically possible solutions. Some of these mathematically possible solutions are not physically possible. There are, however, many physically possible solutions in most cases. These possibilities constitute possibilities for perceptual illusion. Yet the perceptual system commits itself to only one of these solutions, in a wide variety of cases. How is this done?

Many facts complicate this basic problem. The problem has a dynamic dimension. There is a serious theoretical question about the temporal limits of a perceptual

standard of veridical representation. The aim of the psychology of vision is to explain how the perceptual system normally gets things (approximately) right, to the extent that it does, on the basis of sensitivity to light-arrays and other types of input, including internal input.

To solve its paradigmatic problem, perceptual psychology tries to find the "biasing" principles governing how the perceptual system produces perceptual states, type-identified by perceptual representations that are veridical in the cases where they are veridical. The theory should also explain misperceptions, especially where they derive from the normal functioning of the system.

The principles are fitted to the function of the perceptual system in representing (providing perception of!) entities in the environment. The relevant entities are the explanatorily relevant distal antecedents of the proximal light arrays. The theory assumes that perception represents elements in the distal environment. This intuitive assumption is grounded in a larger explanatory scheme. What count as potential perceptual objects—as *relevant* distal antecedents—are roughly those that can be discriminated under appropriate conditions and that are ecologically relevant to the individual's fundamental activities—activities such as eating, navigating, mating, fleeing danger.

I want to give two examples of how empirical psychology postulates biasing principles to explain aspects of visual perception.

Lightness constancy is the capacity to perceive achromatic surfaces (ones that are white, black, or some shade of gray) as having roughly the same surface lightness, despite significant changes in the illumination of the surface or other changes in viewing conditions. The ability to see a page with print on it as having roughly the same shade of white whether one is in a moderately lighted interior or in bright sunlight is an example of lightness constancy. Normal outdoor light is over one hundred times brighter than artificial interior illumination. The amount of light coming off the black print outside is over ten times greater than what comes off the white page inside. Yet normally, we can see the white as the same shade outside as we saw it as inside. Humans and most animals with eyes have lightness constancy. It is one of the simpler of the perceptual constancies.

For various reasons this capacity cannot be accounted for simply by appeal to adaptation of our receptors to changes in lighting conditions, or in any other simple reflexive way. For a variety of reasons that I will not go into, much of the relevant ability centers on responses to ratios of light intensities at luminance edges. A *luminance edge* is a sudden and large discontinuity between adjacent registrations of light intensity by receptors in the retina. Thus if a series of spatially adjacent receptors, which map spatially adjacent light intensities striking the retina, produce a pattern of registrations of significantly different levels of light intensity, the receptors have produced a luminance edge. The roughly constant, or averaged, ratio between sharply different light intensities along relatively local edges is the main starting point for the exercise of lightness constancy. It is known that the visual system can compute such ratios. I will not go into how local ratios are used in forming perceptions of whole scenes. I will focus on only one aspect of lightness constancy.

The light intensity that strikes a receptor is a combination of the reflectance of a surface and the illumination of the surface. So luminance edges are produced by a combination of surface reflectances and surface illumination. The receptors just respond to light intensities. Lightness constancy depends on an ability to separate out surface reflectance from illumination. For lightness constancy is a capacity to track achromatic surface reflectance—brightness of the surface itself—through changes in illumination.

A solution to the visual system's problem of separating surface reflectance (the property that is usually most useful to the animal) from the illumination of the surface is facilitated by the fact that some discontinuities in light intensity are caused by discontinuities in reflectance, whereas others are mainly caused by discontinuities in illumination. Some discontinuities in light intensity in the environment are mainly due to changes in illumination. These are called *illumination edges*. Shadows, reflections on glossy surfaces, differences in surface orientation toward the light source, and focused light sources (such as spotlights) produce illumination edges. Other discontinuities in light intensity are mainly due to changes in the reflectance of the surface. These are called *reflectance edges*. Reflectance edges are patterns of sharp changes in luminance caused by changes in reflectance of two adjacent areas. Much of the problem that is solved by lightness constancy lies in separating reflectance edges from illumination edges. The different distal causes must be separated on the basis of registrations of differences of light intensity that are each combinations of reflectance and illumination.

Key to solving the problem is operating in accord with certain biasing principles. Such principles apply to the registrations of light intensities—in particular, to luminance edges. They specify what perceptions are formed given relevant registrations of light intensities. The principles do not always yield veridical perceptions, but they are fairly reliable in an animal's normal environment. The most basic principles underlying lightness constancy are probably shared by all mammals. I will mention three main principles.

The first centers on the degree of sharpness of the luminance edge. In the absence of information to the contrary, the visual system operates under the principle that a sharp luminance edge is due to a reflectance edge in the environment, rather than to an illumination edge. This principle tends to yield veridical separation of illumination from reflectance, for the most part, because in our actual environment illumination edges—for example, from shadows or spotlights—overwhelmingly tend to be fuzzy, whereas reflectance edges tend to be sharp. The principle can yield illusions under special conditions, however. An extremely sharp-edged spotlight shone onto a surface of uniform reflectance will yield a misperception of a region of higher reflectance on a background of lower reflectance—unless there is some further clue to the presence of the illumination.

A second principle centers on the depth relations among surfaces. The principle states that, in the absence of contrary information, if depth information (which is associated with distance constancy) indicates that two regions are not co-planar, then the edge between the regions is an illumination edge, even if the edge is sharp

Proximal stimulation of the visual system underdetermines the objective facts about a surface. For the image, codified from the geometrical patterns of light emanating from the circles, projected onto the retina is the product of a combination of the actual shapes of the circles and the slant of the page. An irregular texture on a fronto-parallel surface could produce proximal stimulations that are indiscernible at a given time from the stimulations produced by a regular texture on a slanted surface. Thus instead of a pattern of similar, equally distributed circles at a slant, a page that is straight on could exhibit a pattern with nearly regular circles at the bottom, and increasingly flattened, smaller, and more densely distributed shapes at the top. In fact, patterns of this kind on a flat surface seen straight on can produce the illusion of a slanted surface, if there are no other cues available to the perceiver.

The biasing principles produce, however, perceptual representations that favor the slant of a surface with regularly distributed textural elements over the straight-on surface with irregularly distributed textural elements. The biasing principles take as default position that the distribution of textural elements is regular. Given that default position, a projection of light arrays on the retina that derive from foreshortened textural elements combine with the biasing principle to yield a representation as of regular circular textural elements at a surface slant. The biasing principles are mathematically specific, and closely fitted to the empirical evidence about what slants are perceived from what proximal stimulations deriving from what particular distributions of textural elements.

Thus, vision makes use of cues corresponding to the three effects just mentioned. It allows for distortion of *shape* of textural elements, by taking the *foreshortening* of these elements as being in the direction of surface tilt by an amount proportional to the cosine of the slant of the surface relative to the line of sight. It allows for distortion of *size*-image by *scaling* textural elements, so that the relative size of images of textural elements is inversely proportional to the distance of the elements from the eye. It allows for the change of *density* of texture distribution by making transformations on the principle that *an increase in average density* is roughly proportional to an increase in distance.

The biasing principles take cues of the three sorts that indicate that textural elements exhibit a pattern that (roughly) accords statistically with the way textural elements project on to the retinal image for a given surface slant. And they yield a representation of the surface as being at that slant.

The biasing principles depend for their reliability on the pattern of a texture being *regular*. More precisely, textural elements must be *homogeneous*. The statistical relationships among elements on the surface depend only on their relative positions, not on the absolute position in some global reference frame. The statistics of homogeneous planar textures are approximately invariant over translations in the plane of the surface. The notion of homogeneity is more complex for certain curved surfaces. The basic idea should be clear: distributions of shapes, sizes, and densities of textural elements are statistically about the same in any surface region. The patterns that cause illusion are not homogeneous.

registrational and representational states and events are the entities that the psychology describes. The content of the representational states is "explicit" in the system, in the sense that this content marks the representational natures and identities of the main entities (the states and events) postulated by the theory.[15] The principles governing transitions among these states make reference to the states and their contents. Such principles include formulae that are not in any way explicit in the perceptual system—much less attributable to the perceiver.

In every case, biasing principles depend on and mirror basic facts regarding space, motion, light, physical objects, that obtain in the perceptual system's environment.[16] They mirror either laws or deep regularities that hold for the most part.[17] Some principles, including some that govern the simpler perceptual constancies, apply to the visual systems of a wide variety of animals, including some like bees or certain reptiles that surely lack propositional attitudes.[18]

The role of these principles in explanation and in determining kinds of perceptual states brings out some of the fine structure of anti-individualism in perceptual psychology. Perceptual states have representational content that makes reference to the physical environment. The natures of these perceptual states—the perceptual-state kinds—are constitutively dependent on the general character of principles governing their formation. For these principles constitute laws governing transactions among the perceptual states in a perceptual system; and the laws are in a reciprocal constituting relation with the kinds governed by those laws. These psychological kinds and laws reflect and are constitutively determined by kinds and deep regularities or laws in the environment. Not only is the presence of the psychological kinds causally explained in terms of evolution and species-adaptation to the environment. The psychological kinds, marked by their representational contents, are also constitutively dependent on laws, patterns, and kinds in the distal environment of the visual system. So the nature and individuation of perceptual states are constitutively associated, through causal relations, with kinds, patterns, and laws in the physical environment.

I emphasize that the laws governing the formation of perceptual states are laws governing the formation of states *with representational content.* Perceptual-state kinds are type-identified in terms of their representational content, together with the perceptual modality (or intermodality). The representational contents of the states are fixed by the types of transactions into which they enter and by the normal causal and discriminative relations that perceptual states and their associated transformations bear to the physical environment. These discriminative relations are often mediated in conscious beings by the phenomenal aspects of perception. But not all such discrimination need involve phenomenality or consciousness. Many perceptual representations, even in conscious beings, are not available to immediate conscious introspection and may be unconscious in every sense.

There is no getting around the fact that basic kinds in perceptual psychology are intentional or representational. Commitment to representations (and representational contents) as marking perceptual abilities is deeply embedded in the theory's

fail to provide any insight into the intentionality of perceptual states. Either they help themselves to representational power associated with our phenomenology, without explaining that power, or they offer a feature, phenomenology-with-representational-content-bracketed, that does nothing to explicate representational power. Reflection on the role of biasing principles in determining perceptual kinds yields a detailed elaboration of empirical aspects of anti-individualism. Such reflection indicates how perceptual anti-individualism informs and is made specific through empirical explanation.

Empirical psychology does not philosophize about how its fundamental representational kinds are individuated. Its fundamental problem is to explain process not individuation. Nevertheless, its kinds are what they are because of the laws governing perceptual states' getting things right to the degree that they do. These laws are essentially concerned with relations between the individual and aspects of the *distal* environment. In solving its fundamental problem, visual psychology relies on general anti-individualist principles, and fills them in in particular, empirically supported ways. Built into the very methods of visual psychology is the presumption of perceptual anti-individualism. The methods of visual psychology are fundamentally the same as those used in the psychology of other perceptual systems—principally hearing and some aspects of touch.

IV. THE BEARING OF THE EMPIRICAL PSYCHOLOGY OF VISION ON THE INDIVIDUATION OF PERCEPTUAL STATES

I emphasize two aspects of the empirical framework discussed in the previous section.

First, the aim of the theory is to explain the structure of *human* and *animal* perception. The theory does not just explain a mechanism of perception or a set of enabling conditions for perception. The theory does not confine itself to providing an account of a causal chain of non-perceptual processes that precede or lie in the background of an individual's perceiving—and then stop there. The theory incorporates what is known about the accuracy or inaccuracy of whole-organism perception. Perceivers' perceptions and perceptual states are not only the end products of a series of processes. They occur at various stages within these processes. They are often maintained within more complex perceptions that are derived from them. They represent distal environmental conditions that, in veridical cases, originally set the processes going.

Many of the perceptual representations attributed by the theory to *human* perception are known pre-theoretically to be perceptions had by humans. There are perceptions of edges, surfaces, spatial relations, colors, textures, motion, objects—which are familiar to common sense and informed introspection. The methodology of the theory includes tests that make use of this fact. In humans, many of these perceptions, though not all, are accessible to consciousness. Many are easily reported.

imal stimulation, a given type of perceptual state can be produced by different distal conditions. The accuracy or inaccuracy of a perceptual state—including whether a perceptual state is a successful perception of anything in the environment at all—depends on the distal conditions. So the methodology of all serious empirical theory of vision guarantees that given types of visual state can be veridical in some circumstances and non-veridical in others. A perceptual state can be non-veridical in either of two ways. It can be misperception of a particular. Or it can be failure of perceptual reference—failure to perceive any particular.

Failure of perceptual reference is a topic of systematic empirical research. Studies of *apparent motion* are cases in point. In one type of apparent motion, a pair of static images (images of dots, of two-dimensional shapes, or of three-dimensional shapes) are given to the perceptual system at different places and within certain small time intervals. Individuals perceive the situation as involving motion of a single object between the two places. Depending on the stimuli, individuals erroneously experience changes, during the "motion," of shape, size, color, and other properties. With reference to the spatiotemporal interval between the two places, individuals have an illusory perception as of an event, and even plausibly an object, that is simply not there. For example, if red and green dots flash in different places in certain temporal intervals, the individual will have a misperception as of a moving dot changing from red to green. There is no moving object and no event of change of any object's position or color. It is implausible to hold that the subject is misperceiving an actual dot as moving or changing color, since it is unclear which dot it would be. The subject is certainly having a referential illusion of an event of motion and color change.

Apparent motion is studied not only because it provides evidence about how vision works under normal conditions in which slow, meticulous observation is impossible. It is studied also because it has been fruitful in identifying biasing principles. The biasing principles that lead to these illusions govern the perceptual representation of motion and shape under time pressure. In the absence of input that corrects these default biases, motion is perceived as following certain types of paths.[22]

The methodology takes advantage of the fact that a kind of perceptual state can be produced by a given proximal stimulation, whether or not the standard distal antecedents of the proximal stimulation are present. In cases of referential illusion, the biasing principles can be seen to operate under specially controlled circumstances. The scientist can study the biasing principles isolated from the environmental conditions that they normally serve to represent. The content of the representations and the biasing principles are what they are because of reliable patterns of interaction with the environment in the evolution and ontogenetic development of the perceptual system. The theorist can focus on the nature of the representation without distraction from the environmental causes in *particular* cases. Study of other referential illusions has had a similar empirical status.[23]

Perceptual anti-individualism maintains that the nature of a perceptual state type is what it is only because of a *pattern* of normal environmental causes. That

psychology take the individual as a significant unit, and explain perceptual states in terms of impacts on this individual together with internal conditions, especially psychological conditions, in the individual. The negative claim entails that, in *all* psychological explanation of perceptual states, either (a) explanation is non-causal, or (b) individuals and their bodies (hence proximal stimulations) are not relevant units, or (c) psychological causation involves action at a distance, or (d) perception is not of entities in the distal environment. None of these alternatives is attractive or empirically warranted.

V. EMPIRICAL PSYCHOLOGY AND DISJUNCTIVIST FORMS OF ANTI-INDIVIDUALISM

Some forms of perceptual anti-individualism are incompatible with what is known from empirical science. In particular, they are incompatible with the second of the two aspects of the framework of perceptual psychology just laid out.[24] They are incompatible with the Proximality Principle and its consequences for empirical psychology, and indeed common-sense explanation.

The forms I have in mind entail *disjunctivism*. Disjunctivism makes two closely related negative claims. It claims that there is never an explanatorily relevant mental state type in common between (and specific to) a veridical perception and a referential perceptual illusion. And it claims that there is never a mental state type in common between (and specific to) perception of an object and perception of a would-be duplicate substitute for the object that would be, in the context, perceptually indiscernible to the perceiver. The same claims are made with respect to corresponding perceptual beliefs. Disjunctivism makes these claims because it holds that the particular environmental objects (or lack of objects) that are involved in perception are essential to type-identifying all explanatorily relevant perceptual state types and perceptual belief types.

In discussing the duplication case, I will always assume four conditions: (i) that the perceptually relevant *type* of proximal input is the same; (ii) that the antecedent psychological set is the same; (iii) that relevant afferent and efferent internal processes that provide input to the perceptual system are the same; and (iv) that the perceiver cannot in the context perceptually discern through dispositions or phenomenology the difference in substituted objects.[25]

I do not assume, or even believe, that phenomenological indiscernibility suffices for sameness of perceptual state type. I have several reasons for avoiding this assumption. I will mention only one. I think that a phenomenologically indiscernible hallucination, produced by direct stimulation of central areas of the brain, rather than through visual pathways, might well not even count as a perceptual state. This case is ruled out by the assumption that proximal input is the same in the duplication and referential illusion cases. So phenomenological indiscernibility

perceptions and referential illusions are members of different kinds need not be contentious. One can allow that veridical perceptual states make up an important kind (with relevant specific sub-kinds) and that there are non-veridical or illusory perceptual states that make other such kinds. One might think of states and state kinds in a flexible way. There is certainly a referential relation to particulars in all cases of seeing and of veridical visual belief. One might just count being in such relations as states of the individual.[28]

Whether this flexible approach leads to insight is an interesting issue, but it is not our issue. Our issue is whether the disjunctivist denial of common explanatorily relevant kinds of states across the three sorts of cases is correct. It cannot be emphasized too strongly that disjunctivism is not merely the claim that there are mental differences among the three states. In any actual cases, there will trivially be token differences. Some of these token differences, as I shall explain later, are relevant to a certain aspect of representational content. Disjunctivism denies that the relevant perceptual experiences have any explanatorily relevant *type* of perceptual state in common. This is the view that I shall criticize.

Disjunctivism is implausible. Not only common sense but the scientific knowledge that I have outlined support this initial evaluation. Disjunctivism is incompatible with the Proximality Principle, which is basic in nearly all scientific study of perception.

Given that different distal causes can yield proximal stimulation that is relevantly the same, perception of entities in the distal environment is fallible. The Proximality Principle, together with this empirical fact, entails that the same type of perceptual state can be veridical or non-veridical, perceptually referential or non-referential. And the same type of perceptual state can on different occasions be a perception of different, contextually indiscernible particulars. The principle and these entailments are fundamental to the explanatory method and the basic results not simply of particular theories but of the methodology of perceptual psychology itself. Given what we know from empirical psychology, we know that disjunctivism is not a true account of perceptual states.

Disjunctivism entails that token distal differences, in the causal chain leading to perceptual states, that make no relevant difference to proximal stimulation, or to other internal processes that provide input into the perceptual system, or to antecedent psychological states, determine differences in perceptual state types. It further denies that *any* explanatorily relevant perceptual states are caused entirely by these three factors. This view is not only undermined by scientific knowledge. It controverts well-entrenched views about the form of causal explanation in psychology.

Disjunctivism may have been encouraged by running together kind-individuation, causal principles, and particular causal relations on particular occasions. These matters are complex. I want to say just a few things about them here.

There are hypothetical cases, discussed in the twin-earth literature, in which two individuals with similar bodies have relevantly similar proximal stimulation and yet are caused to go into different psychological state types.[29] The key difference between these cases and those required by disjunctivism is that in the twin-earth

It is fairly unusual, at least since the days of Descartes and Newton, for philosophical views to be as directly at odds with scientific knowledge as disjunctivism is. Hegel's claim that there are seven planets comes to mind. It is natural for proponents of disjunctivism and even neutral observers to inquire whether matters might not be so straightforward. Perhaps the science is too young to be counted on. Perhaps the philosophy is making a claim that is not incompatible with the science, because the science is really about a different subject matter than the philosophy.

As to the youth of the science, I believe that any fair-minded comparison of the specific explanatory accounts of vision science with the arguments of disjunctivists will indicate the intellectual and empirical superiority of the former. More importantly, the Proximality Principle that disjunctivism is incompatible with seems necessary to any explanation that would show how perceptions are formed on the basis of proximal stimulation. It is hard to see how there could fail to be perceptual kinds whose individuation is in accord with the principle. The perceptual system can make use only of what is available to it through proximal stimulation and what is present in the background psychology or provides additional *internal* input into that psychology.

It has been claimed that empirical psychology of perception is not about perception but is merely about antecedent enabling conditions of perception. Some claims to this effect are not worth discussing. I will discuss one such claim, intimated by Evans and articulated by McDowell, in the Appendix. I believe that the claim is implausible on its face and has received no serious support. The scientific theory of perception is a theory of the perceptual states and capacities of individuals.

No proponent of disjunctivism has, in a careful and informed manner, confronted the facts from psychology that I have been discussing. No reason has been given to overturn the fundamental modes of explanation and ways of individuating perceptual states and perceptual beliefs that are well entrenched in empirical psychology. Nor has this body of scientific knowledge been shown to be irrelevant to the philosophical issues, or to perception as we ordinarily understand it. The core considerations that inform the science are available to common sense. Disjunctivist forms of anti-individualism have lost touch with fundamentals of what we know about mental states.

The considerations offered in support of disjunctivism are unimpressive. Many depend on phenomenological reflection and strong use of metaphor—appeals to "direct touch with reality" or to what perception "makes manifest." Some rest on programmatic views—such as the view that disjunctivism is needed to rescue us from scepticism. Some are weak arguments in the philosophy of language. Many rest on misunderstanding of what form opposition must take—engaging in criticisms of sense-data views and of an archaic and variously formulated "argument from illusion." I will reserve for the Appendix detailed discussion of prominent arguments in favor of disjunctivism.

I want to remark on one motivation for disjunctivism. The usual motivation is a concern to insure that we make "direct" perceptual contact with the physical

mark perceptual abilities that are fallible. Disjunctivism stymies any natural way of acknowledging the referential fallibility of perceptual states and perceptual beliefs.

In a sense disjunctivism counts the "state" of seeing a given object, say a ball, as referentially infallible. For disjunctivism allows no possible situation in which the "state" type that one is in by virtue of seeing that particular ball could be in error about the existence of the object of perception—that particular ball. Nor could the state make reference to any contextually indiscernible duplicate. For one would not be in the "state" of seeing the ball if there were no referent, or a different referent.

There is a common understanding of seeing according to which seeing does entail a perceived referent. If one sees, there is an object that one sees. Insofar as seeing is further counted a state, it is a referentially infallible state in this sense: no instance of that state could fail to have a perceptual referent. We do not think of our capacity to see as referentially infallible. What we understand by a capacity's being referentially fallible is that it is realized in state types that could have been mistaken. The same state type could have failed to have a referent. This understanding requires that there are perceptual state types that underlie or are involved in cases of seeings, that could have failed to have a referent—lacked a seen object. So even though seeing is commonly understood as factive with respect to the existence of a seen object, seeings involve state types that could have failed to be seeings. They could have been referentially mistaken. These non-factive, referentially fallible state types are the perceptual states recognized by common sense and empirical science.

Disjunctivism denies the existence of such state types. It claims that there are no referentially successful perceptual state types that could have been unsuccessful, and no referentially unsuccessful perceptual states that could have been successful. One can imagine a different account of what it is for perception to be fallible. One could hold that to be fallible is to be indiscernible from a mistaken state. This seems to me, quite obviously, to give an account of a different sort of fallibility—a fallibility in discerning the nature of one's states, not a fallibility in the first-order capacities and states themselves. I think that such a move is clearly philosophically unattractive. For one thing, animals with no reflective powers at all are clearly still fallible in their perceptual states. For another, the view involves a distorted explanation of false (or truth-valueless) but warranted perceptual beliefs. I shall not pursue this matter here. I believe, however, that this problem is symptomatic of the fact that disjunctivism is out of touch with common as well as scientific understanding of human abilities.

VII. VERIDICALITY, PERSPECTIVE, AND ABILITY: ABILITY-GENERAL AND ABILITY-PARTICULAR REPRESENTATION

Perceptual anti-individualism holds that the nature of many representational states, including all perceptual states and perceptual beliefs, is constitutively associated with patterns of causal relations between the environment and individuals. What it

One might think of the same perceptual belief or perceptual representation as being true (or veridical) at one moment, or in one context, and untrue (or inaccurate) at another. I think, however, that the notions of *truth (veridicality) at a time*, *truth (veridicality) in a context* are not basic notions as applied to beliefs and perceptions. I think that beliefs are true or false *simpliciter*. Perceptions are veridical or non-veridical, full stop. To account for truth conditions and correctness conditions, the contents of the beliefs and perceptions in the three cases must be distinguished. There is a singular element in the representational perspective that represents the particular tomato and nothing else.

The singular element in the representational content is required not only by the primacy of truth and perceptual correctness. There are the further considerations mentioned in section I that require including a context-dependent singular element in perceptual content. The practical and representational functions of perception connect the perceiver with particulars in the environment. Representational content must mark—help individuate—the perception of particulars when these functions are fulfilled. Such content must mark the singular representation of particulars in a scene at hand, in distinction from a scene that would be contextually indiscernible to the perceiver.

Particulars cannot be uniquely characterized in perception by context-independent, general representations of properties and relations. Sophisticates can sometimes specify empirical particulars through context-independent means in thought. But they cannot always do so. Normally they do not do so. General attributive categorizations in perception and general attributive concepts in thought are applied to particulars by irreducibly context-dependent means. Context-dependent application to particulars in both perception and perceptual belief has a singular function and is marked by a singular representational content, which can be evaluated for representational success or failure. I label both sorts of representational contents with the term that also labels the acts or events that the representational contents mark and by reference to which these particular kinds of representational contents are individuated. They are *applications*—perceptual applications and applications in thought.

If a person looks at a scene, and an object in it is exchanged with a contextually indiscernible object, the first and then the second object is seen—even though the individual perceiver is unaware of a difference. Different things are seen and different perceptions or perceptual beliefs occur. One belief may be veridical, while the other is not. So different context-dependent singular elements must occur in the representational contents of the perception and the perceptual belief.

Such duplication of scenes is certainly physically possible. Objects at a distance are frequently not seen with such sharpness that a substitute (during a saccade, or as one's attention shifts and then shifts back) would be contextually discernible. Relevant proximal stimulations might even be of the same type as those from the originally seen object. In tracking moving objects, say flies, in a complex scene with distractors, it is quite possible for one contextually indiscernible fly to take the place of another. Perceptual beliefs regarding the two flies might differ in truth value.

The beliefs differ in referential values and truth values, however. The perceptions differ correspondingly, in their reference. Suppose that the first belief is true, the second truth-valueless (or false), and the third false. What differs among these perceptual and belief states is nothing about their general type. They are formed from the same perceptual categories and concepts. They use the same demonstrative singular representational abilities and demonstrative representation types. What is different is a token-dependent difference in the *applications* of the context-dependent singular elements in perception and belief. Thus the different occurrent applications of "that" are marked in the representational content of the belief. Analogous differences in token-dependent singular applications of the perception are marked by context-dependent singular elements in the perceptual content. These are differences in occurrent representation. *Applications* are, paradigmatically, occurrent acts or events which realize singular, context-bound, referential abilities.[32]

Two large categories of representation must be distinguished. The first category is what I have been calling pattern-based representation. This category comprises perceptual and conceptual representational types that mark, or type-identify, general representational abilities, abilities that are *freely repeatable*. These are standing abilities that are not individuated by reference to any *particular* token exercises or applications. Both perceptual representations of kind, property, and relation types and all paradigmatic conceptual representations fall into this category. I call such representations "*ability-general*," or context-free, since they type general psychological abilities—those not individuated by reference to any particular token acts or events.[33]

"Ability-general" is a more precise term than the term I have been using, "pattern-based representation." The relevant representations are not distinguished by whether a pattern is represented. A "pattern-based representation" or ability-general representation can represent a particular, even a particular occurrence. A complete definite description can refer to an occurrence. But the ability underlying the complete definite description need not be constitutively tied to any particular occurrent act, or exercise of a psychological state, or any particular representational event. One can acquire the ability through any of various disjoint sets of events, where the events in the different sets need not even be causally related to one another.

The concept *is orange* is an ability-general representation. One person's mastery of the concept could be individuated by reference to interactions with orange things (and other psychological states) in Africa. Another person could simultaneously learn the concept through interacting with orange things in Canada. There need be no chain linking the two persons' masteries with some ur-event-particular. Thus although the individuating conditions are similar in type, no specific, particular events are essential to their being the same abilities, type-identified by the same representation.[34]

Most concepts and all ability-general *perceptual* representations are general in a further sense. They are capable, according to their content, of applying to various satisfiers of the representational type. They are true of, or veridical of, or applicable to any number of entities. I call such representations "*semantically general*"

text-bound contents that mark applications always combine with ability-general representations in marking representational states. They are, however, a fundamentally different type of representational content. They are individuated in terms of specific occurrent perceptual events.

The representation that marks an application is, I think, best taken to be an abstraction tied for its identity to occurrent events. The representation can mark a memory of a perception. The representation or representational content can mark a mental file, or even an element in an interpersonal chain of thought. Such ability-particular representations are still individuated ultimately in terms of a particular, specific token act (or acts) or event (events). They are not individuated in terms of general abilities or freely repeatable patterns. All instances of such a representation must be individuatively tied by memory or interlocution to some particular act(s) or event(s). The states marked by the representation are thus not *freely* repeatable. Such token-individuated representations in thought are expressed by particular token uses of demonstratives like "that," and by pronomial occurrences taking such tokens as antecedents. As noted, analogs of such demonstrative-applications occur in perception.[35]

Token-individuated singular representation in perception and belief must be differentiated when referents are actually different. Different actual referents require different applications. Similarly, an application with a referent and one without a referent must be distinguished. In the first place, certain ordinary explanations of the particularities of events may use the distinction. These are mostly not the general explanations of perceptual psychology. They are, of course, often equally legitimate. In the second place, semantical evaluations differ. One perception or belief may be veridical (or true) while the other is not. The truth or veridicality conditions must differ.

Such token-individuated context-bound singular representations can differ even when the referent does not vary. As noted, this situation may occur in certain tracking situations or in reidentifications—either in perception or in empirical belief. The difference in application-representation occurs if there is a psychologically relevant, logical possibility that the individual might have been mistaken in taking the object to have remained the same.

Nothing in the singular token-application *specifies* the object, or has a nature that is specific to that object, although the token application is guided by semantically general perceptual or conceptual representational contents. In perception these general, guiding representations never uniquely determine an object in such a way as to guarantee its uniqueness in the world.

Perceptual reference that fails (and produces illusions) can no more be assimilated to purely descriptive, context-free reference than can successful perceptual reference. Yet, clearly, the cognitive context of *different* failures of perceptual reference must be differentiated. The differences have semantical relevance and can bear on explanations of particular events. In this sense, individuation of applications is not dependent on there being a referent.

Recognition of the fallibility of any given perceptual token application supports non-object-dependent individuation of application tokens that in fact succeed in

Third, the primary error of disjunctivism is independent of how to individuate ability-particular, occurrence-based representations. The primary error is the claim that there is no explanatorily relevant perceptual or perceptual-belief *type* common among cases of perceiving an object, perceiving a contextually indiscernible duplicate, and having a perceptually indiscernible referential illusion. There are explanatorily relevant psychological states in common among these three types of cases. The best explanations that we have center on states that are common in their ability-general aspects. Token differences among the states do not affect their sameness of type in this sense. This error is abetted by failure to distinguish the two types of representation. Error could have been avoided by reflection on the Proximality Principle and on psychological explanation.

To summarize: psychological explanation, semantical reflection, and consideration of the fallible and perspectival character of our abilities motivate a distinction between two fundamental kinds of representational content.

One kind marks general abilities. These are the abilities that psychology primarily attempts to provide explanations of. Perceptions that derive from registrations of the same type of light array and other proximal and internal input, given a fixed antecedent psychological set, yield *type*-identical perceptual states and perceptual representations. These identities hold even though the proximal input might be produced by different perceived duplicate objects or by no objects at all.

The other kind marks contextual occurrences. Different instantiations of a perceptual state or belief which have different referents, or which fail to refer, are marked by different representational contents containing different token-individuated representations. These representations mark different applications with different semantical or referential relations to particulars.[38]

Perceptual anti-individualism is fundamentally about the way general representational abilities and states are individuated. Patterns of causal interaction in the individual's or species' history are essential for fixing the nature of types of representational states. When one considers representational elements that are individuated in terms of particular token events of interaction, one must not lump them together with pattern-based individuation of representational types.

Anti-individualistic individuation is present even for token-individuated representations (and hence token-individuated aspects of representational states). If different particulars are *actually* referred to, the token applications that do the referring must be differentiated. If one application *actually* has a referent and one fails to have one, the applications must be differentiated.

Why are applications individuated in a token-oriented way? It is because they mark the individuality of an encounter. The psychological elements in the encounter that are subject to lawful patterns are the elements marked by ability-general representations. From the point of view of understanding patterns of use of available information by perceivers and their perceptual systems, whether an object or its duplicate is perceived on any given occasion is accidental. Similarly, insensitivity to

The key disjunctivist claim entailed by naive realism is negative: No specific explanatorily significant state is common between the different perceptions in the three cases: the case of perceiving object *a*, the case of perceiving contextually indiscernible object *b*, and the case of having a referential perceptual illusion which is for the perceiver phenomenally indiscernible from the two preceding cases.

As indicated, this view is empirically untenable. It is incompatible with a massive amount of empirical evidence central to the scientific study of perception. In fact, disjunctivism is incompatible with the methodology of the science, a methodology that is empirically well supported.

The claim distinctive of naive realism is unacceptable on further grounds. Representational contents are perspectival ways of perceiving. They mark the perspective on objects or properties that the perceiver has in a perceptual state. Direct realism underplays the fact that in human and animal perception, every perceptual reference to any particular object or property—to any perceptual object—is from a partial representational perspective. Veridical perceptual states, like non-veridical ones, have a representational content that marks a perspectival ability, and a way of indicating the physical object or property perceived from a perspective. An account of perception that is true to the perspectival nature of perceptual ability must type-identify perceptions not merely in terms of the objects and properties perceived but in terms of partial perspectival ways of perceiving them. Any physical object or property can be perceived from different perspectives—spatial angles, spatial or temporal frameworks, modes of perception colored by the perceptual apparatus, perspectives that have ego-related implications, perspectives that are unified in perceptual constancies. States with different perspectives on the same object or property are different states.

As argued in section II, representational contents bear a many-one relation to any item perceived. Every visual referent is seen in some perspectival way. The account of perceptual constancy, which I believe is central to our understanding of perception, presupposes the existence of a difference between the way an individual or property is perceptually represented and the individual or property itself. The direct-realist position is not only empirically untenable. I think that it postulates an impossible state of affairs.

The other main view that entails disjunctivism differs from naive realism in maintaining that in all cases of perception of a particular—and perceptual belief successfully about a particular—, there is a representational content that indicates the particular perceived. Perception is perspectival on this view. But perceptions and beliefs are individuated in such a way as to be "object-dependent." Any possible or actual difference in the *particular* that is perceived or perceptually believed-about entails a difference in representational content. This set of views does not yet entail disjunctivism. Disjunctivism follows only when a further claim is added. This is the claim that the representational contents in the different cases help individuate different perceptual or belief state *types*, and that there is no explanatorily relevant

tent. If (2) is false, it does not aid any argument. If (2) is true, it is true only because one cannot know of the existence and structure of perceptual states without background information beyond what is available in current experience. In that case also, (2) does nothing to aid disjunctivism. For the sake of argument, I grant both premises. Yet we should reject the conclusion. We have empirical reason, beyond untutored introspection, to acknowledge representational state types common to veridical and referentially illusional cases.

Snowdon does not provide an argument for disjunctivism. The considerations that he raises either fall far short of the conclusion, or they are tantamount to begging the question.

MCDOWELL

John McDowell propounds a specific version of disjunctivism.[42] On this version, successful perceptual states always involve representational content that refers to the object being perceived. This commitment of the view seems to me correct. Yet, in accord with disjunctivism, the view maintains that any actual or counterfactual difference in the particulars referred to, or any difference between reference and failure of reference, would necessarily involve a difference in perceptual representation *type* and perceptual state *type*—a difference in the *nature* of the state.

In supporting these points, McDowell holds that if an individual's array of perceptual representational states, his perceptual subjectivity, were "exactly as it is however things stood outside it," experience would be "blank"or "dark," instead of "revelatory of the world we live in."[43]

This formulation may appear only to oppose any individualist view that holds that the natures of representational states are completely independent of the environment that is represented. I agree, of course, that such an individualist view cannot make sense of reference to the environment. To be a specific kind of perceptual representational state is to be a state whose nature is essentially associated with types of objects and properties in the environment. Perceptual states could not be what they are if the environment had not been a certain way.

Although the rhetoric may suggest merely this salutary point, the exposition that follows goes much further. McDowell holds that any non-disjunctivist view would leave experience "dark." I can find no argument in the relevant passages that reference would be problematic from the subject's point of view, and that experience would be "dark" or "blank," if an explanatorily significant kind or type of representational state were acknowledged to be common among the different possible perceptual cases. These consequences are not entailed by the denial of disjunctivism. No intervening premises that suffice to entail the conclusion, or even make it more plausible, are stated.[44]

McDowell does give an argument for a different but related conclusion. He claims that as long as it is assumed that the contextual presence of a particular object determines it as the object of a thought and that this presence "cannot enter

Thus McDowell's argument has a significant gap. McDowell gives no reason why a thought or perception's intentional (representational) *nature* must determine the object it is about. He gives no reason why such a nature must fully determine what an individual thinks or perceives. The relevant types of thought and perception seem to determine objects by contextual or indexical elements. These elements' referential success depends on aspects of the context, not on the nature of the representational state alone. Such contextual elements play an ineliminable role in determining what the subject thinks and what the subject's psychological state is. What an individual thinks includes occurrently marked intentional acts, as well as representational types or natures. Representational types or natures guide these acts, but do not fix their reference.

One can individuate representational states for a variety of explanatory purposes. One explanatory purpose relevant to understanding the nature and individuation of representational states is to understand human (and other animal) freely repeatable representational abilities. Some of our perceptual abilities, hence representation types that mark them, can be insensitive to context-dependent, occurrent differences, while still being dependent for their natures on general formative and constitutive relations to the environment. As we have seen, empirical science takes the same view, for more specific explanatory reasons.

Insensitivity to illusion or switch *in special abnormal cases* hardly entails ubiquitous "darkness" or lack of thought with a determinate referential content and a determinate referent. An understanding of the dependence of perception and perceptual knowledge on fallible states, contextual referential devices, and contingent causal relations to the environment is simply part of understanding our limitations. It is part of understanding our dependence on the world for even our surest perceptual contact with it, and our surest perceptual knowledge. There is nothing in the nature or type of a representational perspective that guarantees that our representational kinds or our perceptual-state kinds vary with *these* sorts of occurrent environmental variations.

Nor do anti-individualist arguments that derive from my work demand such a guarantee. In fact, what we know about how perceptual state types are formed from proximal stimulations requires that some perceptual states *not* vary with possible environmental variations *of these sorts.*[47]

In further writing, motivated by commitment to disjunctivism, McDowell tries to cordon off psychology as irrelevant to understanding the nature and form of an individual's perceptual states.[48] I think that this attempt is no more successful than the support offered in favor of the view.

McDowell maintains that there are two compatible but entirely different accounts associated with perception. One is an account of the individual's perception. The other is an account of information processing by the individual's subsystems that makes perception possible. The first is supposed to be that of common sense and Gibson's ecological approach to perception. The second is supposed to be that of the cognitive psychology of perception. The first view attributes "real" representational content. The second's attribution of representational content is

of perception—the structures, kinds, and laws involved in perceiving the physical environment. Their explanations appeal to representational states and contents that are states of both the perceptual system and the individual.

Moreover, Gibson's theory is in empirical competition with representational cognitive psychology, not a separate theory about a different subject matter. The anti-representationalist aspects of Gibson's theory have been empirically defeated. (Cf. note 21.) All sides in psychology agree that the subject matter of visual psychology is the individual's vision. Attribution of psychologically relevant events and representational contents, including some that are attributable only to relatively modular subsystems, are part of an account of what is "constitutively" involved in the individual's seeing objects and properties in the environment. (See the first point in section IV.)

What support does McDowell offer for denying that the empirical psychology of vision explains what is (empirically but "constitutively") involved in the individual's visual perception? What support is offered for the claim that the attribution of representational content to the perceptual system is "irreducibly metaphorical"? Argument is needed. For these positions oppose the intent, practice, and results of empirical psychology. The appeal to representational content is not a metaphor in perceptual psychology. (Cf. section III). Its subject matter is animal and human vision, not merely some set of enabling conditions that are not themselves perceptual states, capacities, or processes. The psychological theory of vision is not a theory of something else besides vision.

Let me quote the crucial part of McDowell's discussion that is supposed to establish his radical theses. This passage purports to support the view that attribution of representational content in empirical psychology is "irreducibly metaphorical":

> What the frog's eyes do for the frog is to put it in touch with moving specks in its spatial environment . . . From the frog's point of view, its eyes enable it simply to pick up the fact that there is a moving speck (with luck, a bug) out there. From the point of view of the frog's 'motor control' (to speak in terms of the 'sub-personal' metaphor), the presence out there of a moving speck is rather (at most) the best hypothesis the . . . whole system . . . can come up with in order to account for the input of light (what is in fact light, though the system does not even know this much) to the eyes. If all goes well, the frog is in direct touch with a feature of its external environment; the internal information-processing system is in direct touch only with structural properties of the immediate inputs to it—which, in the metaphor, it interprets as clues to the nature of the external environment. (Of course the frog does no such thing.)
>
> What could an internal information-processing device really tell an animal? . . . What could an information-processing device *really* tell *anything* (including another component in a sub-personal or 'sub-personal' informational system)? It is essential to realize that the answer to this question can be, in fact is, 'Nothing', without the slightest threat being

environment whose discernment is directly relevant to the animal's basic activities.

Perceptual representations are perspectives on the physical entities that they represent. They mark, or are aspects of kinds of, the individual's perspectival abilities with respect to the environmental entities, types and tokens, that they represent. Neither the registrations nor the representations are themselves objects for the animal or its perceptual system. The perceiver perceives only environmental objects and properties, in ways marked by perceptual representations.

McDowell claims that the frog is in "direct touch" with its environment, but that the internal information-processing system is in direct touch only with structural properties of the immediate inputs to it. I believe that this claim introduces another false contrast.

McDowell does not explain his use of "direct touch." In section V, I explicated two notions of directness—referential directness and non-inferential directness. The referential relations between the perceptual system's representations and the environment do not go by way of representing anything else. So both the animal and the representational states in its perceptual system perceptually refer to the environment "directly" in this sense. Both the perceiver's perception and the perceptual system's perceptual representations are the causal products of a series of transformations that are described by perceptual psychology. The referential "directness" of both the animal's perception and the perceptual system's perceptual representations is dependent on a series of transformations, under principles that have the effect of hypotheses, that convert the proximal stimulus arrays into perceptions.

For the reasons I explained, the transformations are best not counted *inferences,* in the sense that most philosophers use this term. The perceptual system is no less *non-inferentially* direct in its representations of the environment than the animal is. The primary difference between perceiver and system is that the transformations are activities attributable only to the perceptual system, not to the perceiver. The system transforms registrations into perceptions through a series of stages. The animal does no such thing. So the perceiver's perceptual representations—though they are also products of the perceptual system—are not the result of the perceiver's activities, only the results of his perceptual system's transformations. In this weak sense, the perceiver's representations are (relative to the perceiver's activities) transformationally null ("direct"?), whereas (relative to the perceptual system's transactions) they are transformationally complex ("indirect"?).

The claim that the frog is in "direct" perceptual "touch" with the environment, but the subsystem is not, seems to me fundamentally mistaken. Both the frog's perceptions and (what are in relevant cases the same) the perceptions generated by the frog's perceptual system are representationally direct. Moreover, both are generated by a causal chain that runs from distal object to proximal stimulations through transformations by the system.

There is no basis for construing the representational features of cognitive psychology, which are fundamental to its explanations, in a non-realist, metaphorical way. There is no ground for taking the subject matter of perceptual psychology to be different from its declared subject matter—perception.[54]

is no referent, there is no understanding of the use of the demonstrative—presumably also in some "strictest sense." Perceptual beliefs are paradigm cases of demonstrative information-based thoughts. So on his view, there can be no perceptual belief type in common between an instance in which there is a perceptual object and an instance in which there is a contextually indiscernible illusion of a perceptual object.

Evans never explains his "strictest sense." Nor does he explain how this sense connects to psychological explanation. Failures of reference do not obviate the need to explain actions, inferences, beliefs, and desires. No theory of thought that connects to explanation can allow thoughts to disappear merely because of failures of reference. Evans's claims here seem to me idiosyncratic and quaint.[56]

Evans makes two sorts of argument for this form of disjunctivism. One centers on linguistic understanding. Evans maintains (1) that idioms that attribute singular reference in perceptual belief, and perhaps even all ordinary cognitive idioms, "have their home in the activity of interpreting, or making sense of, the speech of others."[57] He then maintains (2) that to understand a singular term intended to invoke identifying information, one must believe that there is something to which the term refers. For, he claims, to justify understanding of an utterance of the form *That G is F*, one must justify a belief that the speaker is referring to some object *a*. He holds (3) that if one does understand such an utterance, one's belief that the speaker is successfully referring to the designated object must be true. For he thinks (4) that understanding an utterance is constituted by knowing some proposition to be true, and he holds that (2) entails that the only such proposition knowledge of whose truth could constitute understanding is the proposition that the speaker successfully refers (to the specific object). He concludes (5) that utterances that contain singular terms intended to invoke identifying information but which fail to secure a referent cannot be understood. From (5) and (1) he infers (6) that no thought involving an indexical or demonstrative singular element can fail to have a referent. Lack of referent entails lack of thought.[58]

This conclusion, together with the view that a particular object helps individuate perceptual belief states, entails disjunctivism.

I believe that every one of this argument's four premises is mistaken—and three of them clearly so. As regards (1): There is perhaps some sense in which cognitive linguistic idioms are first used in understanding other people through their speech. But it is unacceptable to *assume*, as Evans frequently does, that understanding speech is the central issue in giving an account of the nature of perception and perceptual belief. The assumption is out of touch with the empirical study of perception. The account of perception in non-linguistic animals and human infants is distorted if linguistic paradigms dominate the account. Since perceptual belief is dependent on perception, such an assumption will also distort the account of perceptual belief. Even cursory consideration of psychological accounts of these matters shows how far Evans's view is from the natural and normal methodology of the psychology of perception and perceptual belief. The mistake in step (1) is fundamental and renders the rest of the argument largely irrelevant. I shall discuss it anyway.

To understand an application of a demonstrative that fails to have a referent, it is enough to track the historical or perceptual file by relating token demonstrative applications to a single causally connected contextual network. This ability depends on competence with demonstratives and with one's own standpoint. It does not depend on successful reference, or even on belief in successful reference. In the case of understanding others, it requires some appreciation of others' standpoints. One need not share the others' standpoint or perspective.[61, 62]

The foregoing considerations undermine premises (2) and (3). Evans's assumptions (2) and (3) are elaborations of his view that understanding deictic demonstrative uses requires knowing *which object* is referred to. He understood the view as requiring the existence of an object referred to. So understood, the principle is mistaken in light of the sort of considerations just discussed.[63]

The idea that understanding by itself entails or guarantees reference or knowledge seems to me to be correct only in application to certain self-evident propositions. Understanding $2 + 2 = 4$ may yield knowledge of its truth and may entail successful reference to the numbers. Understanding an occurrence of one's own thinking of the *cogito* guarantees knowledge of its truth and a reference for the occurrence of *I*. Understanding ordinary uses of demonstratives in empirical contexts cannot tenably be claimed to entail successful reference.

I turn from premises (2) and (3) to premise (4). Perhaps because it is more abstract, premise (4) is not as obviously false as the other three. I believe contrary to premise (4), that linguistic understanding is not necessarily knowledge that a proposition is true. I believe that inferential capacities and capacities to apply terms correctly in normal conditions are sufficient. I believe that at relative primitive levels of language mastery, a language-understander might have the requisite inferential and applicational capacities, but lack true beliefs that express that understanding. A child need not know propositions that formulate the correct inferences that it can make. All the child's relevant applications of terms may occur in abnormal conditions, yielding perceptual illusions. The child's innate perceptual apparatus could enable it to understand basic empirical vocabulary even though it happens to be brought up in abnormal conditions that lead to false belief. A child need not know the Tarski biconditionals, since understanding does not require a mastery of meta-concepts like *truth,* or concepts of thoughts or sentences. These are large issues.[64] I think that I need not discuss them further here. For the unsoundness of the preceding premises already blocks Evans's argument to the independently unacceptable conclusion that no thought involving a demonstrative singular element can fail to have a referent.

Evans gives one argument for this conclusion that does not rely on a view about linguistic understanding. He claims that we *cannot coherently specify or have a coherent idea of* a demonstrative thought, or an indexical-based thought, absent a referent.[65] This argument avoids the problematic step (1). But it adds nothing to steps (2) and (3) of the earlier argument. It is refuted in the same ways.

As regards (B), I accept that the psychological state is what causes the pattern of use that one makes of the demonstrative. I accept that the way in which one is given the object (a way that marks the psychological state) plays a role in justifying, or warranting, the pattern of use made of the demonstrative. This much seems trivially true. I also accept that justification (or rather entitlement) attaches to the particular belief that is held.

Some qualifications of (B) are needed. The ability-particular or occurrence-based representational element in a particular "way an object is given" is not essential to justifying the pattern of use. Another token could have been part of the same pattern of justification (entitlement) and use, if all else were equal. The representational *kind* or *type* that is part of the way an object is given *is* essential to the justification (entitlement) associated with the pattern of use.[71]

The key step in the argument is (C). The step is not meant to be precise. Insofar as it is supposed to lead to the conclusion of the argument, it is mistaken. An individual can be "justified"—epistemically entitled to his belief—whether or not an object is there. Take an actual case where the believer is entitled to his belief, and the belief is true. Then consider cases of indiscernible referential perceptual illusions or indiscernible substitutions, where the believer has no basis for discerning the difference. In such cases, other things equal, the believer remains entitled to his belief. The believer's being entitled derives from the same entitlement. Perceptual entitlements attach to types of beliefs deriving from types of perceptual belief formation, holding background information and other background circumstances fixed. These types of belief formation derive from *types* of perception formation and conceptualization. Perception formation relevant to warrant is individuated independently of *particular* successes or failures of perception.

Campbell may envisage some other conception of warrant, a factive kind. I believe that such a conception would make a mess of our understanding of warrant, justification, and rationality. What is relevant to the present context is that no argument is given against the straightforward, well-known conception of warrant (or justification) that my view and most "common factor" views rely upon.

The formulations of (C) and (D) suggest the same type-token conflation that has affected so much other disjunctivist reasoning. It does matter to justification (entitlement) whether, as a matter of rule, there is an object there or not. As a type of representational transaction, demonstrative reference is not "indifferent" as between veridical and illusional perceptual states. Justification (entitlement) presupposes a background of reliable, veridical perception in normal conditions. Justification must be a good route to truth. Demonstrative perceptual reference presupposes a background of successful reference and veridical perception. These are matters of rule, pattern, kind, and type.[72]

Even the token act of application is not truly "indifferent" to whether there is an object there or not. Its function is to refer. It refers to the object that causes it. Possible illusions or duplicates are not equally referred to. The identity or "essence" of successful occurrent applications has not, however, been shown to entail the

Campbell's raising them is salutary. I believe that his positive account and his criticism of alternatives are, however, very seriously deficient. I will not present a detailed theory about the matter. I want to discuss three ways of looking at tracking arguments. The first way certainly applies to some cases. The second and third can be refined so as to be compatible. They, or refinements of them, seem to me to apply to other cases.

In some arguments involving demonstratives and tracking, the demonstrative is used anaphorically. In such cases, the validity of the inference can be accounted for straightforwardly in terms of the logical form of the representational content. An individual who maintains anaphorical connections among the demonstrative pronouns can recognize the validity of the argument by reflection on its logical form. The account need not include the object as a "constituent" in experience. Individuals need not anaphorically connect later occurrences of demonstratives to earlier ones in a conscious way. It is, I think, an open question what cognitive and functional conditions result in an individual's connecting different occurrences as instances of a single application. I believe that not all arguments involving demonstratives and tracking fit this model. I believe that some do.

Second, the individual can use separate deictic applications of a demonstrative and assume or "presuppose" identities between the applications. Often such an assumption can be elicited through questioning, or is otherwise plausibly attributable. Sometimes it is not plausible that the connecting identities are represented in the individual's psychology. One can regard the psychology as implicitly relying on what would be a further premise, if it were represented. The further premise-analog would be the analog of an identity statement where different applications purport to be applications to the same particular. Such a premise-analog would provide the argument with an "implicit" validity. If a later application occurs in a perceptual belief about a contextually indiscernible substituted object (or in a failure of perception), the argument would retain its implicit validity, but go unsound.[76]

Refining and specifying the notion of presupposition would be necessary in any refinement of this second conception of demonstrative-involving tracking arguments. There are various notions of implicit representation. And there are probably various types of implicit reliance involved in context-dependent reasoning. These matters are not psychologically simple. In actual tracking cases, it is clear, however, that there is some kind of presumption in the operation of perceptual and belief systems that the perceptually tracked object remains over time, even though one has different perceptions of it—hence, prima facie, different applications to it. Analogs of Fregean informative identities seem to be a part of the natural functioning of perception and perceptual belief. Usually *perception* is deictic not anaphoric. This is shown by the fact that if a contextually-perceptually indiscernible object is substituted for a given object in the course of tracking, the individual will perceive the new object (just as he perceived the old one), even if the individual is completely unaware of the switch. The perceptual reference will usually switch. This

This point about counterfactual situations does not affect the validity of actual arguments involving tracking. The validity of an argument in a particular context is guaranteed by its logical form. The logical form depends on particular occurrent applications as well as general types. No matter whether there is actual failure of reference or switch of reference, and no matter whether other objects could be referred to in counterfactual evaluative situations, an argument is valid if and only if its form is valid in the actual situation. Contrary to (G), what guarantees validity is not experience or successful tracking. Sameness of object cannot guarantee validity. For changes of perceptual perspective on the same object can yield different perception-based logical forms. And experience cannot *guarantee* sameness of object. No sequence of empirical experiences is referentially infallible. In ordinary perceptual tracking, *nothing*—certainly nothing in experience—can *guarantee* "the sameness of object throughout." What guarantees the validity of argument in perceptual tracking is a logical form that maintains the sorts of connection among applications just discussed. So sameness of object in perception is not sufficient for maintaining validity. One must have a valid logical form.

Sameness of object is also not necessary. An object is not necessary at all. One can make valid arguments through referential perceptual illusions, using anaphora. By assuming or presuming mistaken identifications through perceptual tracking, one maintains validity without soundness. Again, it is the logical form that matters, not the presence or sameness of the object of perceptual reference. Moreover, validity may not be necessary for being warranted.

Campbell offers a few pages of criticism of my version of the "common factor" view. Discussion of this criticism may be too detailed for some readers. They may happily skip to the next section. In the interests of explicitness, however, I will respond to the criticism.

The main idea of Campbell's criticism derives from Evans's conception of understanding discussed earlier. The idea is that the "common factor" view cannot account for our *understanding* of perceptual beliefs about particular objects. Campbell extends this point by claiming that the view cannot account for our perceptual knowledge of particular objects. I believe that Campbell does not articulate notions of understanding or knowledge fully enough to have given a genuine argument. The key shortfall is that no good objection is given to the idea that we can understand, and have knowledge through, a perceptual belief even though the content of the belief, in other circumstances, could have made reference to substitutes or could have been involved in illusion.[78] I lead up to Campbell's main claims, and my replies, gradually.

Campbell correctly characterizes my view as attributing a demonstrative element in the content of perception and perceptual belief. He does not distinguish the demonstrative form from the occurrent application. At this stage in the argument the distinction is not crucial.

I quote at some length his discussion of my view. I published the views that he discusses many years ago, but the relevant features of those views are unchanged. He holds that my account

erentially unsuccessful. Like other non-veridical representation, non-veridical experience depends (in the ultimate individuation of the representational state) on veridical, *de re* experience.[82] Nothing in Campbell's discussion shows why *de re* perception and *de re* perceptual belief turn out, on this account, to be any less fundamental than they seem to be.

So much for general points. Let us turn to Campbell's more specific charges.

(J) and (K): The perceiver's "subjective life" includes (a) a demonstrative capacity (marked by a schematic representation analogous to *this*); (b) an occurrent application of this "demonstrative element" in a context; and (c) the perceptual attributives, in many cases with their associated phenomenologies, that mark repeatable perceptual abilities and that characterize purported aspects of purportedly perceived particulars. The aspects of the context that, together with this three-element perceptual representation, fix a perceptual referent are not available to the subject *in this sense*: To perceive a particular, the individual perceiver need not be able to describe or otherwise represent as such the causal connections that fix the particular that is perceived. This is a piece of common sense.

The perceptual occurrence that constitutes the application is "subjectively available" in the sense that it is part of the exercise of the perceiver's abilities. Self-conscious believers can be aware that they are doing this. On this understanding, I accept (J) and (K). I think that no dire consequences follow.

(L): The "subjectively available" application within perception does, in cases of successful perceptual reference, distinguish the particular entity perceived in the context. It picks out that particular and not some other one that would be perceptually indiscernible if it had been substituted. Neither the application-governed perception nor the perceiver could discriminate the particular from a *would-be* duplicate that *might have been* substituted (or that might come later), or from a *would-be* phenomenologically indiscriminable perception that *might* fail to refer. That is a matter of empirical fact. In the actual context of successful perception, the perceiver and his perception distinguish or discriminate the particular by perceiving *it* and not the would-be substitutes. Perceiving it involves applying the attributives in a singular way in response to the effect of that object. The particular referent is not discriminated from all possible or actual look-alikes through general abilities that apply across contexts. It is discriminated ineliminably in a context-dependent way.

This account matches the perceiver's abilities. Apart from perceiving a particular, the perceiver cannot discriminate what he actually perceives from would-be duplicates or illusions. This inability does not diminish the fact that the perceiver and his perceptual representation represents the particular that is perceived. The lack of would-be discriminability is exactly what one should expect. The account of content should match the account of ability, since content marks ability.

Campbell does not explain "of itself" in (L). The demonstrative- (application-) governed perception does discriminate the perceived object from all other objects. Of course, the application depends on its association with general perceptual

exists and is not a contextually indiscernible duplicate. There is no defense of this assumption. It would amount to begging the question. Nearly all serious views of knowledge would reject the assumption. It is enough for knowledge that the possibilities of a switch or an illusion are not contextually relevant threats and that one has no reason to think that such threats are in play. Assurance lies in the reliability of experience, the way the nature of perceptual representational *kinds* depends on the nature of the environment, and so on. One need not know these facts in order to have warrant grounded in them. No guarantee against error is necessary.

(P): It is true that "meta" knowledge of the context is not normally available. It is also true that specific knowledge that could discriminate the context from contexts where reference fails or where duplicates are substituted is not normally available. Such knowledge is not needed for understanding or for perceptual knowledge.

MARTIN

Michael Martin offers a complex argument for the naive realist version of disjunctivism and against what he calls the "intentionalist theory."[85] The argument begins with a sub-argument for what he calls "The Dependency Thesis":

(DT) to imagine sensorily a phi is to imagine experiencing a phi,

or to take the specialization of the thesis that he works with:

(DT') to visualize a phi is to imagine visually experiencing a phi.

The considerations pro and con regarding these theses are interesting and complicated. I do not find the theses convincing, however. It seems to me that in a certain clear and natural sense, one can visualize an object and *not* imagine visually experiencing the object. One imagines the object from the perspective of a visual experience, but no experiencing of the object (either by oneself or by anyone else) is imagined to be included in the imagined scene. One might visually imagine El Capitan in Yosemite from a certain perspective, but not imagine any experiences of it. One might even imagine that there are no experiences of El Capitan in the imagined scene. Prima facie, this seems to be visually imagining El Capitan without imagining an experience of El Capitan. Prima facie, there seems to be nothing contradictory in this claim, as there would be if (DT') were (constitutively) true.

Whether or not there is a sense in which (DT') is true, Martin's argument for it certainly fails. He begins by rightly noting that visualizing an object involves taking an imagined visual perspective on the object—for example, visualizing it from a perspective according to which the object is to the left. It is quite true that one could have such a perspective on the object only if one were to have an experience of the object. It does not follow that if one imagines something from a perspective that one *could have* only if such and such were the case (only if one were experiencing the object from that perspective, or only if there were an experience of the object from that perspective), then in imagining something from that perspective one *must imagine* such and such to be the case. Prima facie, the experience-like visual imaginings that are employed in imagining El Capitan need not themselves be the objects of visual imagination as such.

text of the *transparency* of visualizing (the lack of an introspectively evident medium) is difficult for the "intentionalist" to account for.

As he notes, the obvious thing for his opponent to claim is that "the state of imagining itself comes with a commitment to the imagined situation's being a certain way." This should be understood as a reassertion of the fact that visualizing entails placing the visualized object in the imagined scene.

Then, oddly, Martin replies that this claim "fails to take into account the full consequences of the Dependency Thesis":

> Certainly, in imagining a visual experience one is thereby actually committed to there being a visual experience in the imagined scene, but (A') the extra move that is needed is a commitment to the presence of what the imagined experience is an experience of. When one entertains the supposition that there is a pig in the room, one does not have to entertain the supposition that one believes that there is a pig in the room. (B') What the intentional theory is required to do is to explain how in imagining an experience with a certain content one thereby also takes up a similar suppositional attitude toward the content of the imagined experience.[87]

I believe that there is a confusion here about what needs to be explained. (A') The commitment to the presence of what the imagined experience is an experience of is not an extra move. It is part of the intuitive starting point. We are granting (DT')—that visualizing a pig is imagining an experience of a pig. We are also granting that visualizing a pig requires visualizing a scene in which there is a pig. So we are granting that in visualizing a pig, one is imagining a *veridical* experience of a pig, an experience that is factive. Visualizing a pig is placing a pig in the imagined scene. So any experience that one imagines as part of the visualizing must itself be taken to be veridical with respect to the imagined scene.

(B') The "intentional theory" is not required to explain how by merely imagining an experience with a certain representational content one *thereby* supposes that the content is veridical. I deny that this supposed *explanandum* is even true. In the imagined experience which, according to the Dependency Thesis, is involved in visualizing, one does not take the representational content of the imagined experience itself to involve or entail the presence of the visualized object. Rather one has begun with the supposition of veridicality. One simply takes the content of the imagined experience to be veridical. The content itself might not have been veridical. That is the position of the "intentional theory." Certainly, it is my position. To assume otherwise would be to beg the question.[88]

Martin writes, "There seems to be a serious challenge here for the intentional theorist to explain the phenomenology of our sensory imagining, to explain how such imagery seems to give us the presence of an imagined scene rather than a mere imagined experience of the scene."[89] This claim again seems to evince confusion about what is to be explained, or else to beg the question. I hold that the imagery does not by itself guarantee the presence of the imagined scene rather than a mere

naive realism or of disjunctivism. The points are natural expressions of the central claim that Martin is contesting. Neither the perceptual experience in itself nor the imagery in "factive" visualization in itself ("in itself" in the sense: given their metaphysical identities) "include" the objects, or entail their existence. Martin's argument for a naive realist form of disjunctivism has no force.

****

I believe that the various disjunctivists have collectively failed to make a single successful point in favor of disjunctivism—or against the natural alternative. The fact that the theory is incompatible with empirical knowledge should give pause to those tempted by the arguments, assertions, and metaphors that mark this line of thought.

NOTES

I have learned from comments by Ned Block and Johannes Burge.

1. Many philosophers distinguish representations and representational contents. Representations are token, structured instances of, or vehicles (like inner symbols) expressing, representational contents. Representational contents are abstractions that help type psychological kinds. Although this distinction is plausible in application to perceptual psychology, I am non-committal about whether it has a useful general application in psychology. The need for shareable, repeatable representational contents is very clear in both perceptual and cognitive psychology. The need for psychological states with representational contents is also very clear. In this paper and elsewhere, I use "representation" and "representational content" interchangeably, in effect leaving out representations in the sense just outlined, except insofar as they are the states themselves. In any case, I shall avoid the ontological issues that token representations raise.
2. I think that anti-individualism may not apply to certain aspects of thoughts about the qualitative features of sensations. I doubt that the thought *This is pain* depends on relations to the environment beyond the body for its specification of pain. By contrast, the aspect of the mental state marked by the present-tense index does, I think, constitutively depend on such relations.
3. For an elaboration of this framework in a context that centers more on the functions and norms governing simple perceptual systems, see my "Perceptual Entitlement," *Philosophy and Phenomenological Research* 67 (2003): 503–48, especially section I.
4. For a much more detailed account of different types of general elements in perception, and their various roles, see my "Five Theses on *De Re* States and Attitudes," in a volume honoring David Kaplan (New York: Oxford University Press, forthcoming). Also see section VII.
5. I believe that singular representation suffuses the perceptual representation. I think that every general element (property-, kind-, or relation representation) is accompanied by a singular element that represents a particular that purportedly instantiates the property or relation attributed. Such particulars are in addition to particulars that purportedly *have* the properties or stand in the relations. I do not rely on this belief here. I assume only that perception normally represents particular objects, events, and places and attributes properties, kinds, and relations to them.
6. The neglect stems, I believe, from a combination of a correct view—that perception is non-conceptual/propositional—with an incorrect view—that singular reference requires a backing of concepts. Cf. P. F. Strawson, *The Bounds of Sense* (London: Routlege, 1989), originally published 1966; cf., e.g., 97–112. An example of neglect of singular elements in perception is Gareth Evans, *The Varieties of Reference*, ed. John McDowell (Oxford: Oxford University Press, 1982). Evans develops a notion of non-conceptual content that carefully avoids allowing singular elements in the content (cf. 122–25, 181). Evans's view influenced subsequent writers—for example,

bodily motion. Gibson's insensitivity to the importance of explaining illusion is now widely recognized in psychology. Cf. Roger N. Shepard, "Ecological Constraints on Internal Representation: Resonant Kinematics of Perceiving, Imagining, Thinking, and Dreaming," 417–47.

12. Most of the points made in the last eight paragraphs can be found in introductions to any mainstream work in visual psychology. For example, see Stephen E. Palmer, *Vision Science* (Cambridge, MA: MIT Press, 2002), 9–11, 18–24, 55–59, 247–48.

13. An account of lightness constancy can be found in Palmer, *Vision Science*, 122–33. For more background, see E. H. Land and J. J. McCann, "Lightness and Retinex Theory," *American Journal of the Optical Society of America* 61 (1971): 1–11; Irving Rock, *The Logic of Perception* (Cambridge, MA: MIT Press, 1983), 279; A. L. Gilchrist, "Lightness Contrast and Failures of Constancy: A Common Explanation," *Perception and Psychophysics* 43 (1988): 415–24; A. L. Gilchrist, *Seeing in Black and White* (Cambridge, MA: MIT Press, forthcoming); A. L. Gilchrist, "Perceived Lightness Depends on Perceived Spatial Arrangement," *Science* 195 (1977): 185–87; E. H. Adelson, "Lightness Perception and Lightness Illusions," in *The New Cognitive Neurosciences*, ed. M. Gazzaniga (Cambridge, MA: MIT Press, 2000).

14. For detailed discussion, see David C. Knill, "Surface Orientation from Texture: Ideal Observers, Generic Observers, and the Information Content of Texture Cues," *Vision Research* 38 (1998): 1655–82; "Discrimination of Planar Surface Slant from Texture: Human and Ideal Observers Compared," *Vision Research* 38 (1998): 1683–1711. For elementary discussion, see Stephen E. Palmer, *Vision Science*, 234–36. I do not know whether the textural elements are registrations of pre-perceptual retinal information or early perceptual representations of two-dimensional shapes on a surface. Here, the answer does not matter.

15. This is a minimum characterization of the "explicit" presence of perceptual states and their representational contents. I believe that there are causally relevant token realizations—neural or psychological—of each component in an "explicitly represented" perceptual content in the perceiver. I will not go into this complex matter here.

16. For an article that discusses other biasing principles in an illuminating way, see Elizabeth Spelke, "Principles of Object Perception," *Cognitive Science* 14 (1990): 29–56.

17. For a discussion, see Roger N. Shepard, "Ecological Constraints on Internal Representation: Resonant Kinematics of Perceiving, Imagining, Thinking, and Dreaming."

18. C. R. Gallistel, "Animal Cognition: The Representation of Space, Time, and Number," *Annual Review of Psychology* 40 (1989), 155–89; Charles R. Gallistel, *The Organization of Learning* (Cambridge, MA: MIT Press, 1990); C. R. Gallistel, "Insect Navigation: Brains as Symbol-Processing Organs," in *Invitation to Cognitive Science*, vol. IV (Cambridge, MA: MIT Press, 1996); Christa Neumeyer, "Comparative Aspects of Color Constancy," in *Perceptual Constancy*, ed. Vincent Walsh and Janusz Kulikowski (Cambridge, UK: Cambridge University Press, 1998); David Ingle, "Perceptual Constancies in Lower Vertebrates," in *Perceptual Constancy*, ed. Vincent Walsh and Janusz Kulikowski.

19. There is a conception of objective, relational, "perspectival" properties that is popular among some philosophers. This notion goes back at least to C. D. Broad, *Scientific Thought* (London: Routledge and Kegan Paul, 1923), chapter 8; *The Mind and Its Place in Nature* (London: Routledge and Kegan Paul, 1925), chapter 4. It got its more recent impetus from Gilbert Harman, "The Intrinsic Quality of Experience," in *Philosophical Perspectives*, J. Tomberlin, ed., vol. 4 (Atascadero: Ridgeview Pub. Co., 1990). These properties are commonly taken to be phenomenologically accessible. For example, perspectival size is taken to be a plane of a size that would exactly occlude vision of an object from a certain viewing distance. I do not deny the existence of such properties. The association of those like size and shape with phenomenological aspects of experience, however, seems to me a mistake. There are effectively infinitely many plane sizes set at different distances between any given object and any given viewing position that would occlude the object. None is objectively privileged. It is not plausible that all of these sizes are phenomenologically accessible as such. Phenomenology cannot be explained purely by reference to the individual's relations to objective, relational, non-psychological properties. I believe that it is a philosophical and scientific mistake to regard any objective "perspectival" properties, such as perspectival size, shape, color, as among the objective environmental entities seen, unaided by background theory. Cf. note 8. I hold a similar position on the views of Sydney Shoemaker in "Introspection and

Nilsson, *Perspectives on Memory Research: Essays in Honor of Uppsala University's 500th Anniversary* (Hillsdale, NJ: Lawrence Erlbaum, 1979), the journal *Ecological Psychology*, and various books on ecological psychology. Although this work contributes insights, it provides no convincing ground for resisting several decades of experimentally supported, mathematically sophisticated explanations that appeal to representation in visual perception. This aspect of Gibsonianism is a fringe position.

There are a few near-adherents to Gibsonian ideas in philosophy of psychology as well. Cf. several essays in *The Embodied Mind*, ed. F. J. Varela, E. Thompson and E. Rosch (Cambridge, MA: MIT Press, 1991); J. K. O'Regan and Alva Noë, "A Sensorimotor Account of Vision and Visual Consciousness," *Behavioral and Brain Sciences* 24 (2001): 939–73; and Alva Noë, *Action in Perception* (Cambridge, MA: MIT Press, 2004). Some authors in this group make, or intimate, wholesale criticisms of representationalist theories. They provide no detailed empirical criticism or detailed empirical alternatives. The criticisms are offhand in comparison to the science being criticized. Often this body of work presents the mainstream representationalist approach in misleading and sometimes mistaken ways. (A side remark: the empirical phenomenon, change-blindness, presented in the 2001 paper as an original motivation for opposing representationalist theories, in fact seems to support standard representationalist theory, although change-blindness is in one of the less-well-developed areas of vision science. Cf. D. J. Simons and R. A. Rensink, "Change Blindness: Past, Present, and Future," *Trends in Cognitive Science* 9 [2005]: 16–20.) The science has gone too far to be vulnerable to empirically unspecific, wholesale anti-representationalist criticism of this sort. Some representational aspects of vision are very well understood. There are literally hundreds of empirically well-supported, mathematically detailed explanations of aspects of vision in the representationalist framework. If one requires less in the way of mathematical detail, there are thousands. These explanations can be overturned only through demonstrations of specific empirical errors or through presentations of comparably detailed, superior empirical explanations.

22. Visual representation depends in some cases on input from non-visual modalities—such as the vestibular system or proprioceptive signals from the eye and head. In representing simple matters like motion of a physical body under time pressure, it depends little or none on thought. Cf. Z. W. Pylyshyn, "Is Vision Continuous with Cognition? The Case for Cognitive Impenetrability of Visual Perception," *Behavioral and Brain Sciences* 22 (1999): 341–65, esp. 361.

23. Cf. Shepard, "Ecological Constraints on Internal Representation," 422: "Precisely because our own internal constraints so well match the external constraints in our world, these internalized constraints reveal themselves only when externally available information is degraded or eliminated. Being tightly controlled from without, activity in the perceptual system is then necessarily guided more by whatever constraints operate within." Cf. also Shimon Ullman, *The Interpretation of Visual Motion* (Cambridge, MA: MIT Press, 1979); Roger N. Shepard, "The Role of Transformations in Spatial Cognition," in *Spatial Cognition: Brain Bases and Development*, ed. J. Stiles-Davis, M. Kritchevsky, U. Bellugi (Hillsdale, NJ: Lawrence Erlbaum, 1988); Roger N. Shepard, "Perceptual-Cognitive Universals as Reflections of the World," *Behavioral and Brain Sciences* 24 (2001): 581–601. This mode of study carries over to the study of imagery and the rotation experiments for which Shepard is famous. The content of images is widely seen to derive from perception. Perceptual illusions have loomed large in showing that perception is governed by constraints that mirror regularities of the physical world. Cf. Irving Rock, *The Logic of Perception* (Cambridge, MA: MIT Press, 1983), 319; J. A. Wilson and J. O. Robinson, "The Impossibly-Twisted Pulfrich Pendulum," *Perception* 15 (1986): 503–4; A. L. Yuille and S. Ullman, "Computational Theories of Low-Level Vision," in *Visual Cognition and Action: An Invitation to Cognitive Science*, vol. II, ed. Daniel N. Osherson, Stephen M. Kosslyn, and John M. Hollerbach (Cambridge, MA: MIT Press, 1990); Stephen E. Palmer, *Vision Science, passim*, e.g. 365.

24. I think that most supporters of anti-individualism reject disjunctivism. I will not cite specific works. The interested reader would do well to consult work by Ned Block, Fred Dretske, Donald Davidson, Jerry Fodor, Barry Stroud, Michael Tye, and others. Their work is not centered on disjunctivism, but much of it involves anti-disjunctivist commitments.

25. Because I am assuming that perceptually relevant types of proximal stimulation are the same, cases of change-blindness are not at issue. In those cases, the scenes are not contextually indiscernible to perceivers. Perceivers just commonly fail to discern the differences.

34. The notion of ability-generality is also not the contrary of object-dependence-on-a-particular. If there is object-dependent representation that represents a concrete particular and that is also learnable by any number of different people who could have no relation to one another, and who are not related to any ur-act (e.g., a baptismal act), then the representation would count as ability-general, even though it is individuated partly by reference to the object of reference.

    Although concepts expressed by numerals are, I think, individuated partly by reference to their referents, and are thus object-dependent, such concepts are ability-general. Their individuation hinges purely on types of psychological activity and psychological capacity—not on any particular, specific, occurrent acts or exercises of capacities.

35. I discuss this sort of singular context-dependent representation, insofar as it occurs in thought, in "Belief *De Re*"; "Russell's Problem and Intentional Identity," in *Agent, Language, and the Structure of the World*, ed. James Tomberlin (Indianapolis: Hackett Publishing Company, 1983); "Vision and Intentional Content," in *John Searle and His Critics*, ed. E. Lepore and R. Van Gulick (Oxford: Basil Blackwell, 1991). I discuss such singular representations insofar as they occur in perception and thought in "Perceptual Entitlement"; in "Descartes and Anti-Individualism: Reply to Normore," *Reflections and Replies: Essays on the Philosophy of Tyler Burge*; and at greatest length in "Five Theses on *De Re* States and Attitudes." Cf. also "Postscript: 'Belief *De Re*,'" in my *Foundations of Mind: Essays by Tyler Burge*, Volume II (Oxford: Oxford University Press, 2007).

36. Our evaluations of modality are two-sided in this regard. We can abstract from the referent and truth value of a perception-based thought and conclude that the same thought event, with the same perspectival content—with its token-dependent elements—might, in a different world, have failed to refer. Then the associated thought might have had a different truth value. The type-individuated representations are the same, and we can imagine that the same application token-event had a different immediate-contextual-causal ancestry. Then the thought event that is in fact true could have been false or truth-valueless. This point undergirds our view that representations in perceptual beliefs—both concepts and applications—are referentially fallible in any particular instance. An analogous point applies to perceptual representations. Nevertheless, when we evaluate the modal status of a thought considered as a truth (or a perception considered as veridical), we take the singular applications in perceptual beliefs to be rigid. We evaluate the modal status of the actually *true* thought even in situations in which it might have been false. That truth might have been false if the object had not had the relevant properties. We would not be evaluating the same truth if in considering counterfactual situations, we take the object of reference out from under it. Some perception-based thoughts can thus be regarded as necessary truths—necessarily: that object is self-identical; necessarily: that object is a fruit.

37. I leave open the full individuation conditions of applications: for example, whether applications can be the same if they are guided by different fundamental ability-general, semantically general attributives.

38. Token applications are also crucial in accounting for the connection of perception and perceptual beliefs with particular acts—such as initiating ingestion of food or attack on prey. In this paper I focus narrowly on the role of perception in the *acquisition* of information.

39. From a broader perspective, one can see reference to any particular individual as the result of lawful interactions. Still, psychological explanation succeeds because it centers on psychological kinds that are not sensitive to all referential illusions and duplications.

40. All serious defenders of disjunctivism that I know of are British. In addition to insular influences, this is to be explained, I think, in terms of a preoccupation with the excesses of the British empiricists, and of their British successors in the early and early-mid part of the twentieth century. The idea of disjunctivism is suggested briefly by J. L. Austin, *Sense and Sensibilia* (Oxford: Clarendon Press, 1962). The first defense of disjunctivism that I know of is J. M. Hinton, "Visual Experiences," *Mind* 76 (1967): 217–27. Hinton's motivations are not clear, but I think that rejection of sense-data is probably central. This seems to be the motivation of Paul Snowdon, "Perception, Vision, and Causation," *Proceedings of the Aristotelian Society* 81 (1980–1981): 175–92; and "The Objects of Perceptual Experience," *Proceedings of the Aristotelian Society*, suppl. vol. 64 (1990): 121–66; and Michael Martin, "The Reality of Appearances," in *Thought and Ontology*, ed. M. Sainsbury (Milan: Franco Angeli, 1997). Some of the arguments in John Campbell, *Reference and Consciousness* (Oxford: Clarendon Press, 2002), have this motivation.

51. It does not follow, nor is it true, that genuine perceptions (as opposed to registrations of information) are conscious. There is massive evidence that a good bit of human perception—which involves constancies and a rich representation of the physical environment—is unconscious.

52. "The Content of Perceptual Experience," 197–98. McDowell goes on to say that animals are "semantic engines" and to indicate his understanding of the role the instrumentalist use of content attribution plays in psychological explanation:

> We could not make sense of the competence that enables us to make sensible use of the claim that *animals* have dealings with content if we could find nothing inside them but, say, a completely homogenous jelly. And nobody knows how to make sense of an animal's internal control mechanism, and connect it conceptually to the competence it is supposed to explain, except by describing it *as if* it were, what we know it is not really, a semantic engine, interpreting inputs as signs of environmental facts and, as output, directing behaviour so as to be suitable to those facts in the light of the animal's needs or goals. To insist that the attribution of content at this sub-personal or 'sub-personal' level is 'as if' talk is in no way to debunk it … it is surely clear, at least in a general way, how content-attribution that is only 'as if' can even so pull its weight in addressing a genuine explanatory need: the question is what enables us animals to be the semantic engines we are.

It is certainly not true that no one knows how to describe the animal's modular psychology without describing it as if it interpreted signs. Attributions of *interpretation* to the subsystem, or indeed the animal, play no role in the theory.

53. Perhaps McDowell is assuming that if an animal or subsystem knows nothing it cannot have genuine (non-metaphorical) representational content. Such an assumption would need argument. It is out of keeping with the practice of empirical psychology. Perceptual psychology attributes perceptions to lower animals that lack knowledge. The perceptions have objective reference to objects and properties in the environment and exhibit perceptual constancies much like ours.

54. There are many signs that McDowell thinks that the veil-of-ideas picture is endemic to cognitive psychology's explanations. In addition to the passage just quoted, cf. "The Content of Perceptual Experience," 197, 200, 203. He seems to think that the modular system is "blocked off" from the environment, and that regarding the system's representational states as non-metaphorical would commit one to taking the animal to have "direct" access only to its "interior." Such views misunderstand the science. I think that they are encouraged by taking the telling metaphor and (even worse) the interpretation metaphor as literal parts of the science. Some of the difficulty may lie in McDowell's accepting Dennett's remark about the brain's being a syntactic engine (which I regard as already at best misleading) and transferring it, without discussion, to the perceptual system. McDowell describes the perceptual system as a physical mechanism, ibid., 198. That may be so in a metaphysical sense. From the standpoint of science, it is basically a *psychological* system. Theorizing about syntax in psychology is vastly less developed and central than theorizing about changes among states with representational content. Syntax (the form of perceptual organization and processing) serves representational content. It has no serious role in visual psychology independent of it.

55. McDowell's appeal to Gibson seems to me ill-considered. Gibson saw his work as empirical explanation of seeing of objects on the causal basis of sensitivity to ambient light. He did not found a science of ecology separate in its objectives from empirical psychology. McDowell's "reading" of Gibson as doing something completely different from what cognitive psychologists do is not borne out by how Gibson's work has been assimilated and criticized in psychology—as any number of textbooks will bear out. Gibson's denial of a role for representational states in the causal account is empirically discredited on numerous grounds. The status of Gibson's work ultimately is, of course, not the central issue. (Cf. note 21.)

56. Gareth Evans, *The Varieties of Reference*, 136. Evans's strange view that no thought is thought when there is a failure of reference is inessential to disjunctivism. His holding this view shows how far his approach is from reflecting seriously on psychological explanation or representational abilities. Although Evans made important contributions in directing attention to issues about thought, I think that he was much too fixed on considerations regarding language and semantics in his

ibid., 411. This statement is true—at least if (DT') is true. However, it is not strong enough from the point of view of his opponent, who distinguishes sharply between having a visual experience as of an apple which is veridical and having a visual experience as of an apple that is illusory. The cases are very different even if the type of representational content is the same. The cases are different even if counterfactually or in the imagination, the very same type *and* token could have remained the same. In visualizing an apple, one imagines a *veridical* visual experience as of an apple—if in visualizing one imagines a visual experience at all. I believe that Martin allows himself to interpret the Dependency Thesis in a way that takes his opponent to regard the imagined experience as neutral as to whether it is veridical, and then challenges his opponent to come up with an explanation of how the veridicality is to be accounted for. This mistake is also suggested by Martin's handling of an objection at ibid., 412–13, and in the main body of his argument against the "intentionalist," which I am in the midst of discussing.

89. Ibid., 416.

90. Ibid., 417–18.

91. Such placement is subject only to very general constraints. One cannot visualize something in the scene both to have a property and to lack it, from the same perspective—although one can perhaps visualize Escher cases which present impossibilities, incompatibilities that emerge from different perspectives on the same image. One cannot visualize something and its absence in the same visualization. The object must have visually discernible properties. And so on. The point is not that visualization is in all ways infallible. It is that its sources of fallibility are more limited than and different from those of visual perception.