

Perceptual Entitlement*

TYLER BURGE

University of California, Los Angeles

The paper develops a conception of epistemic warrant as applied to perceptual belief, called "entitlement", that does not require the warranted individual to be capable of understanding the warrant. The conception is situated within an account of animal perception and unsophisticated perceptual belief. It characterizes entitlement as fulfillment of an epistemic norm that is apriori associated with a certain representational function that can be known apriori to be a function of perception. The paper connects anti-individualism, a thesis about the nature of mental states, and perceptual entitlement. It presents an argument that explains the objectivity and validity of perceptual entitlement partly in terms of the nature of perceptual states—hence the nature of perceptual beliefs, which are constitutively associated with perceptual states. The paper discusses ways that an individual can be entitled to perceptual belief through its connection to perception, and ways that entitlement to perceptual belief can be undermined.

Epistemology after Descartes tends toward hyper-intellectualization. It also tends toward misplacement of cognitive differences between mature humans and other animals. Many philosophers have exaggerated how much we understand and control our cognitive lives. Many have proposed denying non-human animals cognitive capacities that there is overwhelming empirical reason to believe that they in fact have, on the ground that such capacities require having language, self-consciousness, ability to understand reasons, or the like. It has been common to analyze knowledge and warrant in terms of doing one's epistemic duty, or being epistemically responsible, or being able to access justifications in a space of reasons, or following accessible rules for the direction of the mind—abilities only mature humans seem to have. It has also been common to deny that animals have beliefs about the physical world, and even to deny that animals have genuine perceptual systems, as

* I am indebted to Tony Anderson, Paul Boghossian, Martin Davies, Richard Foley, Elizabeth Fricker, Mikkel Gerken, Peter Godfrey-Smith, Gilbert Harman, Paul Horwich, John MacFarlane, David Papineau, Christopher Peacocke, Teresa Robertson, Stephen Schiffer for comments that led to changes. I benefited from giving earlier versions at the Epistemology Conference at Rutgers in the Spring of 2000, and subsequently at Brown University, New York University, the University of Kansas, the University of London (the Jacobson lecture), the University of California, Berkeley (one of the Townsend lectures), and at a conference on my work at the Australian National University, Canberra, summer of 2002.

opposed to other sensory systems. As a consequence, human cognitive systems have been understood in abstraction from their animal roots. I believe that all these tendencies are instances of hyper-intellectualization.

This tradition has been challenged in recent years by philosophers advocating forms of epistemic externalism. Epistemic externalism is true of some types of warranted belief and some types of knowledge. It constitutes an antidote to hyper-intellectualization. However, the view commonly carries an overt or covert message that traditional conceptions of epistemic norms—of reasons and justification—are irrelevant to understanding knowledge or cognition. Often justification is replaced by, or analyzed in terms of, natural association of ideas, reliable processes, or law-like transmission of information. On many such radicalizations of epistemic externalism, epistemology is to be replaced by psychological, sociological, or historical studies of how human beings actually operate.

This message compromises the core insights of epistemic externalism by attaching it to an ideology that loses touch with the role of epistemic norms for perceptual belief and with the role of norms of reason and justification in many higher types of knowledge. It combats hyper-intellectualization with an even less attractive anti-intellectualism. Thus, thermometers and connectionist machines replace God-like critical reasoners as paradigms for human cognition.

I will explicate an epistemically externalist sub-species of epistemic warrant, which I call "*entitlement*". Entitlement is *epistemically externalist* inasmuch as it is warrant that need not be fully conceptually accessible, even on reflection, to the warranted individual. The individual need not have the concepts necessary to think the propositional content that formulates the warrant.¹ Entitlements that I shall discuss are epistemically externalist in the

¹ I introduced the notion of entitlement in "Content Preservation", *Philosophical Review* 103 (1993) pp. 457-488; reprinted in *Content*, Villanueva ed. Ridgeview Press Atascadero, Ca. 1995; and in *Apriori Knowledge*, Casullo ed., Hampshire, Ashgate Publishers, 1999. I developed it in another context in "Our Entitlement to Self-Knowledge," *Proceedings of the Aristotelian Society* (1995), pp. 1-26; reprinted in *Externalism and Self-Knowledge*, Ludlow and Martin eds., CSLI Stanford University, distributed by Cambridge University Press, 1998. In the intervening years, the term "entitlement" has become popular, used sometimes my way but sometimes rather indiscriminately to express a wide variety of other epistemic notions—including even epistemically internalist notions of warrant. In my own work, I insist on the term's being understood in the way I originally introduced it—as expressing a notion of warrant that is epistemically externalist inasmuch as it does not require that the warranted individual have in his repertoire concepts needed to formulate or think the propositional warrant for the belief. Hence the notion contrasts with the internalist form of warrant—justification. Individuals can be epistemically entitled to a belief without having reasons warranting the belief, without having the conceptual repertoire necessary to have relevant reasons for the belief, and without having the concepts needed to understand or even think the entitlement. I think that any being that has entitlements to beliefs must have justifications for some other beliefs—or at least must be capable of having justificatory support for some other beliefs. For I think that having

further sense that the warranting features include relations between the individual and an environment. The other primary sub-species of epistemic warrant is *justification*. Justification is warrant by reason that is conceptually accessible on reflection to the warranted individual.

I want to integrate the epistemically externalist notion of entitlement with *anti-individualism*—often called “externalism” in the philosophy of mind. This is the principle that certain representational states are dependent for their natures, and for the individuation of what representational contents they have, on certain relations between the individual—or relevant representational systems of the individual—and certain aspects of the environment that is represented. I believe that perceptual entitlement is partly grounded in the anti-individualist character of perceptual states and perceptual belief.²

I shall operate under a background assumption that reasons and justification are significant factors in some types of knowledge and are more sophisticated sub-species of the same epistemic genus—warrant. Here I concentrate on one of the simplest forms of entitlement—perceptual entitlement.

A viable conception of warrant and knowledge must include both primitive and sophisticated types. A condition of viability is that such a conception apply to animal and child perceptual belief, and knowledge, as well as to mature human instances of belief and knowledge. I regard it as empirically well established that animals have representational perceptual systems, and that higher non-human animals have systems of belief.³

I.

Entitlement is a subspecies of epistemic warrant. Epistemic warrant, and hence entitlement, is an epistemic good. The *epistemic* good, warrant, is

beliefs requires being able to carry out inferences, and the relevant inferences must sometimes support beliefs. Inferentially supported beliefs are justified, or are at least backed by beliefs that are capable of justifying them. Higher non-linguistic animals have reasons or justifications for some of their beliefs in this sense. My notion of having a reason or a justification does not require reflection or understanding. That is a further matter.

² Tyler Burge, “Individualism and the Mental,” *Midwest Studies in Philosophy IV* (1979), pp. 73-121; “Other Bodies,” in *Thought and Object*, A. Woodfield, ed. (London: Oxford University Press, 1982); “Cartesian Error and the Objectivity of Perception,” in *Subject, Thought, and Context*, McDowell and Pettit, eds. (New York: Oxford University Press, 1986); “Individualism and Psychology,” *The Philosophical Review* 95 (1986), pp. 3-45; “Intellectual Norms and Foundations of Mind,” *The Journal of Philosophy* 83 (1986), pp. 697-720.

³ Cf. e.g. Noam Chomsky, Review of *Verbal Behavior*, by B.F. Skinner, *Language* 35 (1959) pp. 26-58; Charles R. Gallistel, *The Organization of Learning* (Cambridge, MA, MIT Press, 1990); Richard Byrne, *The Thinking Ape: Evolutionary Origins of Intelligence* (Oxford, Oxford University Press, 1995); Michael Tomasello, and Josep Call, *Primate Cognition* ((New York, Oxford University Press, 1997); Marc D. Hauser, *Wild Minds* (New York, Henry Holt and Company, 2000). I believe that higher animals have simple sorts of perceptual knowledge.

essentially associated with the fundamental *representational* good—truth. The notion of an epistemic good must be understood in relation to this fundamental representational good. Although truth is fundamentally a representational good, it is also an epistemic good, because belief, warrant, and knowledge are apriori connected with truth.

The representational function of a system of beliefs is to form true beliefs. What true beliefs it is important for the system to form varies with the system, and depends on other functions of the system. But the representational function of any such system is to represent veridically—to yield true beliefs. Given that an animal has a psychological system with its attendant representational functions and representational goods, there will necessarily be subsidiary functions associated with norms for representing as well as possible given the perspectival, informational, and structural limitations of the system. The epistemic goods *warrant* and *entitlement* are fulfillments of norms associated with achieving the representational good, truth. Warrant, and its sub-types, are not the only epistemic goods. But warrant is a conceptually central one—worthy of the philosophical interest that it has attracted.

Truth is a subspecies of *veridicality*. Truth is propositional. Other sorts of veridicality include those of a photograph, phonograph recording, map, and so on. Non-propositional types of representation are like propositional representations in having conditions of veridicality. The veridicality of non-propositional representations, like photographs, maps, perceptions, will be termed “*correctness*”.

Truth applies to representational *propositional* contents regardless of what individual considers the content. I use “*representation*” and “*representational content*” interchangeably. I am non-committal about the relation between representations, which I regard as abstract, and token instances in the mind or brain. I am more intensely non-committal on whether representations are instantiated or expressed by tokens that have an independent “syntactic” or neural specification.

The notion of warrant that interests me applies not primarily to representations or representational contents independently of some representer. It applies to individuals’ representational faculties, abilities, commitments, practices, states and events.

Whether an epistemic state is warranted depends not only on whether it provides, in a certain sense, a good route to truth. The notion of epistemic warrant is also conceptually tied to considerations regarding limitations of perspective, information, and structure in the particular system or animal at hand. Whether an epistemic state is warranted depends partly on the epistemic perspective of the individual in that state—what information the individual has or what information is accessible, and how well that information is used. Epistemic states can be warranted even though their representational content

is not veridical. Being warranted entails being well-positioned to achieve or indicate truth in normal circumstances, given certain limitations on the individual's abilities and perspective.

I will not try to say, precisely and in general, what limitations count. Broadly speaking, the limitations derive from "accidents" of the individual's positioning in the world, or other ignorance that does not stem from defect or misuse of the individual's intentional systems. Thus warrant is compatible with *brute error*—error that does not reflect badly on the functioning or use of the individual's representational systems, relative to the function of representing veridically.⁴

Warrant is also compatible with "brute truth"—truth that is accidental, relative to the good use and functioning of the individual's intentional systems. Brute truth is the stuff of Gettier cases. The individual has true warranted belief, but lacks knowledge. This is because the warrant does not connect to truth in the right way. The individual is "lucky". How to explain the relevant senses of "luck" or "accident" is another difficult problem which I shall not try to solve.

Being warranted entails being reliable, at least under certain conditions. The reliability necessary for being warranted must be associated with the nature and function of the individual's representational systems, and with the way representational content is established for these systems. The reliability must obtain in standard use in normal circumstances. I shall later make "standard use" and "normal circumstances" more precise. One does not become warranted because a bolt of lightning or some puppeteer sets up a reliable connection to a subject matter. One does not become unwarranted by landing in an abnormal environment, or being systematically tricked by a demon or scientist.

The account that I am giving is frankly teleological. I assume that there are certain functions, ends, goals, and commitments, which bring with them goods for animals and their subsystems. Goods are relative to these ends, goals, commitments, or functions. I use "good" in a very broad sense to apply to a fulfillment of a function, end, goal, or commitment, or to the fulfillment of some standard or norm associated with furthering or fulfilling functions, ends, goals, or commitments.⁵

⁴ Cf. my "Individualism and Self-Knowledge" *The Journal of Philosophy* 85 (1988), pp. 649-663, for a discussion of brute error.

⁵ I leave open whether there are goods with an "all things considered" status, and whether there are intrinsic goods. In evaluating an animal, one may want to ask what it is good for it to do or be, taking into account all its functions—practical, representational, and so forth. I do not rule out there being answers to such questions. Or one may ask whether certain goods are good not simply in being a fulfillment of a function (or the like), but in being in some sense more absolutely or intrinsically good. I do not rule out there being positive answers to this sort of question either, although I believe that such answers would probably be the product of showing that the functions which such goods fulfill have a

All ends, goals, and commitments, and many functions, are those of the whole animal. But the well-functioning of a sub-system can be a good for the sub-system, when it fulfills a function of the sub-system. A heart's fulfilling its function of pumping blood efficiently is a good for the heart relative to that function. It fulfills that function of the heart. Fulfilling that function might also be a good for the whole animal, inasmuch as it contributes to fulfilling some function, end, goal, or commitment that the animal has. This might be the function of mating, or the goal of mating with a particular animal.

Many ends, goals, and functions can be established as such on biological grounds. I take it that survival, at least long enough to have offspring or to fulfill some other biologically basic functions, is an end for all animals—an end that can be established biologically. Human beings can have other sorts of ends. An end brings with it subsidiary goals and functions. Fulfillment of these goals and functions are goods for the animal, relative to those goals and functions, and ultimately relative to its end or ends. Thus finding food at a given time is a fulfillment of a deer's goal of eating on a given occasion, hence a good for the animal, relative to that goal. It is also a fulfillment of the deer's biological function of eating. Ultimately, finding food is a good relative to the deer's survival-oriented end. Similarly, for the well-functioning of the animal's subsystems. Fulfillment of the function of a subsystem, like the coronary subsystem, is a good for the subsystem, relative to that function.

There are distinctively psychological functions. The visual system of many animals has the psychological function of representing colors and shapes, and perhaps of representing things as predators or as food. I think that these psychological functions are not reducible to biological functions, at least if biological functions are taken to be fundamentally non-representational. But psychological functions are closely associated with biological functions. As with most biological functions, psychological functions such as those of representing shapes or predators are established empirically.

Some functions and their attendant goods can be established *a priori*, given the presence of certain psychological capacities. For example, a good for a perceptual system is to represent veridically, indeed reliably. I take it as obvious that it is known *a priori* that the central *representational* function of a perceptual system is to perceive. This function is *a priori* associated with a representational function (to represent veridically). This representational function is also empirically associated with other functions or ends, such as the

fundamental or universal status in certain sorts of valuation. This is, I suspect, why it has been traditionally thought that truth, moral goodness, and perhaps beauty, have some sort of "intrinsic" value. But none of these issues is germane to the present discussion.

function of veridical perception in helping the animal find food, to contribute to navigation by the animal, to contribute to its biological end, and so on.⁶

I think it *a priori* that a *representational* function of a psychological system of belief is to form true beliefs. Understanding what a belief is suffices *a priori* to warrant the view that such a system has a function to represent truths. Beliefs aim at truth, and are defective in a certain way if they are not true. Fulfillment of such a function is a good for the system. A system of

⁶ After developing the broadly Aristotelian framework that I am about to elaborate, I found some points of contact with Ernest Sosa, "Intellectual Virtue in Perspective" in *Knowledge in Perspective: Selected Essays in Epistemology* (Cambridge University Press, Cambridge, 1991). I avoid the terminology of "virtue" epistemology because it suggests intention and control (although Sosa eschews these suggestions). In what follows, I emphasize norms in ways he does not. Sosa's account is also relativized in various ways—for example to an environment—that mine is not. This leads to various differences in judgment about what is warranted. In my view, the environment which helps ground the content of the representations is *the* relevant environment for considering warrant. I argue this point in section VI. As a further consequence, there is no counterpart in his system to my notion of entitlement. In particular, his notion of aptness is different. On the other hand, the broad teleological cast and the lack of reductive ambition of my account are similar to his. Fred Dretske, "Entitlement: Epistemic Rights Without Epistemic Duties?" *Philosophy and Phenomenological Research* 60 (2000), pp. 591-606, uses essentially my notion of entitlement, and develops a general picture that I find largely congenial. One difference is that, as with Sosa, he tends to consider reliability primarily in relation to the environment the entitled individual *happens* to be in (correctly arguing that so construed reliability is not necessary for entitlement). An emphasis on function is also found in Ruth Millikan, *Language, Thought, and Other Biological Categories* (Cambridge, Mass., MIT Press, 1984); and Alvin Plantinga, *Warrant: The Current Debate* (New York, Oxford University Press, 1993), and *Warrant and Proper Function* (New York, Oxford University Press, 1993). My view was developed independently, even though later. I do not share Millikan's reductionism. I do not think that psychological much less epistemological notions are fundamentally biological, even though their use is in certain respects influenced by broadly biological considerations. I do not take the notion of function to be more basic than the epistemic notions that I connect with it. My usage of "warrant", and my concept of warrant, are to be distinguished from Plantinga's. In Plantinga's terminology, warrant is whatever distinguishes knowledge from true belief. On his usage, warrant is undermined by Gettier cases. My usage—which allows for true warranted belief that is not knowledge—is, I think, more standard. There are thus cases in which a state is warranted in my sense, but not in his. There are cases in which a state is warranted in my sense, but is not even an instance of proper functioning (relative to the function of producing true beliefs) in his. Cf. note 24. Still, there is a broad congeniality between some aspects of the two approaches.

The section VI deduction, in something like Kant's sense, of perceptual entitlement, with its use of anti-individualism and its specific emphasis on the role of the normal environment (in a sense to be developed below), is distinctive to the present account, as is the integration of empirical work on perception. Moreover, the emphasis on specifically representational and epistemic functions associated with psychological states, and the view that the representational functions of the perceptual system and belief system are knowable *a priori* is not, as far as I know, to be found in the other accounts cited in this note.

Incidentally, I think that the *a priori* principle that the representational function of a perceptual system is to perceive, hence represent veridically, is more substantive and informative than the logical truth that a perceptual system is a perceptual system—because of the role of the concept *function* in the principle.

belief formation has other functions—to guide behavior, for example, or perhaps to contribute to survival. It is crucially important in what follows to distinguish representational and epistemic functions from practical functions and from more primitive biological functions. A given psychological system commonly has various functions.

The notion of function that I appeal to in discussing representational functions is primitive and intuitive. The notion of a biological function has often been regarded as “explained” in evolutionary terms. I do not expect the notion of a representational function to receive any analogous “explanation”. Still, it is a notion that we understand and should make use of. Whether it is essentially associated with some notion of design is itself an open question for me. I take the notion to be essentially connected to our basic concepts of intentionality or representation. I take these to be, at least *prima facie*, primitive.

There are those who ignore or attempt to explain away the representational functions of perceptual and belief-forming systems. They see biological or practical functions as the only relevant ones. The good of belief is judged purely relative to such functions. I think such views clearly mistaken. There is no prospect of eliminating the notions of perception and belief from our understanding and explanation of ourselves and other animals. Given an objective use for the notions of perception and belief, and given the *a priori* arguments in the second and third paragraphs back, I think it clear that there are fundamental representational functions associated with psychological phenomena. It is in principle open to investigation whether these notions of representational functions can be reduced to other notions, for example, practical notions. But I think that such reductions are philosophers’ pipe-dreams. I need not argue this here. What is important here is that representational goods and functions have an objective standing, not that that standing be non-derivative.

There are certainly deep non-accidental relations between the practical functions of belief and its representational function. True belief helps in survival, in finding mates, and in other practical ways. The representational functions and contents of sense-perceptual systems are inseparable from biological and practical functions of animals that have those systems. For biological functions of the whole animal play a role in determining the representational functions of perceptual systems and in individuating perceptual contents. Hence biological and practical functions figure in individuating many beliefs. On the other hand, representational goods and functions are distinct from practical goods and functions. Being true is not in general being useful. It does not follow that representational functions are not in some way reducible to practical functions. But I believe that reductive enterprises here tend to be more exercises in ideology than contributions to knowledge and understanding.

Of course, it is an empirical question, to be answered by psychological methods, whether some given system *is* a perceptual system or system of beliefs. But if a psychological system is a perceptual system, it is apriori that its fundamental and characteristic function is to perceive, hence to represent veridically. Analogously, for a psychological system of beliefs. Such a system may have other functions. Its characteristic function is its representational function: to provide true propositional representations. It is apriori that these systems have representational functions and that these functions are associated with goods in my broad sense.

The relation between being a good for a subsystem and being a good for the animal is sometimes a delicate matter. It must always be borne in mind that goods are goods (fulfillments) relative to functions, ends, commitments, goals. The functions of a subsystem are not necessarily functions of the animal. So though fulfillment of a function of a subsystem is commonly a good, relative to some goal, commitment, function, or end of the animal, the connection is not in general an apriori one.

As background for the predominantly epistemic discussion, I shall be assuming *perceptual anti-individualism*. This is the view that the individuation and nature of perceptual states and perceptual beliefs are necessarily associated with relations between the perceiver and aspects of the environment. The nature of a perceptual state, as marked by its representational content, is partly determined by relations between the perceptual system, perhaps in its evolutionary history, and features of the environment that cause instances of its states and that states of the system represent.⁷

⁷ Cf. my "Individualism and Psychology" *The Philosophical Review* 95 (1986), pp. 3-45; "Cartesian Error and the Objectivity of Perception", *Subject, Thought, and Context*, McDowell and Pettit, eds. (New York: Oxford University Press, 1986) 117-136; reprinted in *Contents of Thought*, Grimm and Merrill, eds. (Tucson, Arizona Press, 1987); "Descartes and Anti-Individualism: Reply to Normore" forthcoming in *Reflections and Replies: Essays on the Philosophy of Tyler Burge*, Hahn and Ramberg, eds. (Cambridge, Mass., MIT Press, 2003). It should be noticed that perceptual anti-individualism does not imply that each perceptual state is what it is because instances have been caused by features in the environment that it represents. It is possible for there to be perceptual state types that have never represented veridically. A perceptual state may misrepresent something as having a particular shape; and no instance of that perceptual state has ever been caused by light coming from anything with that (represented) shape. What perceptual anti-individualism demands is that the perceptual state's representational content be part of a system of shape-representing content which is grounded in some relations between shapes in the environment and instances of some perceptual state-types, relations which help explain representational success.

The grounding relation does not require that there be a type of state that is specifiable independently of its representational content that was brought to have its representational content through interaction with the environment. Things do turn out that way in some instances. What the grounding relation requires is that wherever perceptual states have perceptual content, this content must be associated with a background of functionally significant causal relations between elements in the environment and that state (or appropriately related representational states)—and those elements are what the state rep-

The anti-individualist framework enriches the point that perceptual systems, and perceptual states, have the representational function of representing veridically, hence reliably. For natural applications of anti-individualism explain the presence of the perceptual systems and perceptual states in an animal as a causal consequence of their enabling the animal to successfully represent the environment. Being capable of representing successfully in favorable, natural circumstances is a consequence of the system and states' being present in an animal. These two facts instantiate a general scheme for functionalist explanation that is quite common in biological explanations of a system's function.⁸

Representational goods are fulfillments of representational functions. Representational goods are necessarily associated with *norms* for fulfilling such functions. Like the notion of a good, the notion of a norm is to be taken very broadly. A norm need not be understood or intentionally adhered to by the individual that it applies to or governs. Norms need not guide. Given the

resents. This is a metaphysical necessity. The anti-individualist relations are explanatory in that they help show how representation is possible. Similar sorts of relations can be expected to play a role in a genetic explanation of the perceptual states, but the form of that genetic explanation may or may not attribute a prior state-type (for example a neural state type) that comes to have the relevant representational content. And the individual-environment relations postulated by such a genetic explanation will normally be metaphysically contingent.

Throughout this essay, I emphasize the role of evolutionary origins in the determination of perceptual representations, especially perceptual categories. This is because I think that the role of innate mechanisms in determining basic perceptual principles and basic perceptual categories is in fact huge. This emphasis is not, however, to be taken as suggesting that I think that it is apriori that evolution played this role, or as suggesting that there are not many *learned* representations that are applied in perception. Learned representations occur not only in high-level perception (where perception is supplemented by concepts of objects that cannot be categorized in purely perceptual terms—for example, concepts that rely on notions of function or natural kind), but even in simple perceptual groupings by the perceptual system that make use of innate perceptual categorizations. It is just that I think that the role of bottom-up, innate perceptual mechanisms has been underappreciated by many who are not tuned to the empirical literature. For overview examples of the literature that are relevant to these points, see Z.W. Pylyshyn, "Is Vision Continuous with Cognition? The Case for Cognitive Impenetrability of Visual Perception" *Behavioral and Brain Sciences* 22 (1999), pp. 341-365, and Roger N. Shepard, "Ecological Constraints on Internal Representation: Resonant Kinematics of Perceiving, Imagining, Thinking, and Dreaming", *Psychological Review* 91 (1984), pp. 417-447.

⁸ See Larry Wright, "Functions," *Philosophical Review* 82 (1973), pp. 139-168. Wright presents his account as an analysis of the meaning of "function". I do not accept this view. But I do think that he provides a close approximation to a correct account of a widespread way of connecting functions to causal explanation in evolutionary biology. I do not accept the idea that representational functions are, in any ordinary sense, biological functions. The notion of function that I employ is not to be explicated in terms of this pattern of explanation, or even taken to be materially equivalent with any notion of biological function. Cf. note 7. I suggest only that noting the connection between this pattern of explanation in biology and anti-individualism in psychology and philosophy of mind might be illuminating.

psychological functions and capacities of an individual or psychological subsystem, there are standards for achieving a good. So goods generate shoulds. Relative to a good, there are normative standards for the well-functioning of the system, or for the exercise of the individual's abilities. Epistemic norms are certain types of standards governing fulfillment of representational functions of psychological states. The standards that we are interested in govern how the individual or system should operate, given its representational functions and given the limitations on its information and capacities for representing well. Epistemic norms govern the acquisition, transformation, and employment of *representational contents* in a system of belief. They are standards for the formation of certain representational states, processes, and methods in fulfilling the representational function of representing veridically. They are standards for representing reliably and well, given the capacities and perspective of the individual or subsystem. They are norms governing achievement of the representational good, truth.

Epistemic warrant is an epistemic good that consists in meeting such epistemic norms. Thus epistemic warrants are goods in that they fulfill certain norms that are associated with fulfilling the representational function of certain representational states. Epistemic warrant attaches specifically to representational systems, representational methods, representational competencies, or representational states—or to the individual's reliance on them. It does not attach to stimulations, or to registering of pre-representational information, or to just any neural processes that might be associated with psychological systems. It also does not attach to representational systems that do not have psychological states, such as computers or thermometers, though other norms governing the achievement of veridicality certainly do apply to artifacts whose representational capacities are parasitic on those of their users or programmers.

I re-emphasize that the applicability of epistemic norms does not presuppose individual control. Such norms need not even be comprehensible to the individual whose activities fall under them. Animals need not be able to understand or be guided by the norms that in fact apply to their formation of beliefs. On the other hand, the norms do not rest purely on reliability. They apply when and only when certain psychological competencies are present. Given the representational abilities of the animal, and given its perspectival limitations, there are standards that govern states and transitions that well serve the function of forming true beliefs. The relevant epistemic norms come from this group.

The fact that the norms are external to an individual's knowledge and control does not make their applicability any less real. Animals have real psychologies. The relevant functions, ends, and goods are apriori associated with the psychology of autonomous representational states. Epistemic goods entail

standards or norms for furthering truth, given the individual's psychological capacities. The norms that govern an individual's psychology depend on abilities and limitations of the individual.

The idea that in attributing representational systems, and their functions, goods, and norms, we are merely taking up some stance, or the idea that the attributions should otherwise be viewed non-realistically, has no more plausibility or force here than it does in any other domain in which attributions are integral to an empirically or intellectually well-founded endeavor. Psychology is such an endeavor, as is the everyday attribution of propositional attitudes. The psychology of vision, which I shall be centering upon, is an exceptionally advanced, and scientifically sophisticated area of psychology. I shall not argue the point. I will simply assume that the psychology and epistemology being discussed are to be understood realistically, for non-human animals as well as for people.

II.

I presuppose a distinction between *perceptual systems* and other sub-systems that have functions and success conditions for interaction with the environment. Template matching in the production of genetic material, phototaxis in protozoa, and detection of oxygen in fish have some things in common with perception. They involve "getting things right" relative to some function or end. But they are not perceptual systems. These systems involve only a degenerate "getting things right". An interesting sub-case of this point is the distinction between perceptual systems and mere (non-perceptual) *sensory systems*. The distinction is drawn in psychology and biology, as well as by common sense. I am content here to let empirical theorizing draw the line, a line which will inevitably betray fuzzy edges.

As a first approximation, where invocation of computations on representations, which carries an attendant intensionality in the theorizing language, is a source of genuine empirical insight; where the representations are available to the functions of the whole animal; and where the system shows perceptual constancy—an ability of sub-systems to filter out the effects of perspective and other extraneous factors in homing in on a system-independent subject matter—, there is reason to take the attribution of a perceptual system seriously.

Template matching systems and detection of oxygen in fish appear to fail the first, and hence the third, of these conditions. They also fail the second: The fish visually perceives its food; the detection of oxygen is not something attributed to the whole animal; it is something—as is digestion—that its subsystems alone do. Phototaxis and sensing of heat by protozoa fail the first

and the third conditions.⁹ Vision is a prototypically perceptual system. Vision is to be distinguished from phototaxis.

On the phylogenetic scale, genuinely perceptual systems appear in animal species well before belief and propositional attitudes appear. I call animals that lack propositional attitudes “*lower animals*”. Animals that have propositional attitudes are, correspondingly, *higher animals*. Bees, frogs, pigeons, goldfish, and octopi are, I assume, lower animals. Although they lack propositional attitudes, they have visual perceptual systems. The perceptual systems of some of these animals have been thoroughly studied. Scientific explanations of the discriminations, computations, and informational functions of the perceptual systems of lower animals commonly individuate the representational content of visual states partly in terms of properties and relations in the animals’ environment, properties and relations to which the animals bear causal relations—both in sensory reception and in activity. In fact, best explanations of some of these low-level representational systems attribute perceptions of physical objects in space, and of rudimentary spatial features of and among such objects. For example, computations in the visual system of bees that bear on locating a hive operate on parameters that represent spatial positions and objects in those positions.¹⁰

Lower animals exhibit perceptual constancies and discriminatory behavior that make it clear that their visual systems are representing features of the physical world, not simply reacting to stimulus frequencies, and not simply carrying information in the sense that tree-rings carry information. Honeybees, goldfish, and pigeons, as well as primates, exhibit color constancy—the ability to distinguish a given color despite radical differences in illumination that would affect the light intensities falling on the eye. Most lower animals show similar constancies in shape-, distance-, and motion perception. Lower animals are able, for example, to compensate for rigid rotations in visual rec-

⁹ Fuller discussion of perceptual constancy, of what can be attributed to the whole animal, as opposed to its sub-systems, and of the distinction between sense-perceptual systems and sensory systems that are not perceptual systems, must be postponed to other occasions.

¹⁰ There is considerable empirical evidence that lower animals have what might be called categorial rather than merely dispositional perceptual representations. Many such representations apply to objective properties like shape, spatial relations, even physical bodies, and are not individually tagged in relation to the animal’s needs or point of view. Some in the Gibsonian tradition used to claim that animals perceive only “affordances”—properties definable only in terms of the animal’s needs or dispositions, such as *danger* or *shelter*. It seems empirically clear that this view is incorrect. Insofar as there are perceptual representations of *danger* they are essentially associated with perceptual representations of types of shape, movement, and so on. They are commonly formed (perhaps innately) through grouping representations of properties that signal properties relevant to the animal’s needs. On another occasion, I will critically discuss philosophical positions that have attempted to forestall attribution of perceptions of objective properties to lower animals, and indeed many higher animals.

ognition, and can hold real-world rather than body-centered coordinates of some objects in memory. The navigational systems of insects show sophisticated computational capacities in mapping the space around them.¹¹

If one reflects on the role of perceptual systems in lower animals' finding something to eat and in avoiding being eaten, and if one sheds behaviorist or rationalistic prejudices, it should not be surprising that the representational content of their perceptual systems specify objects and properties in the physical world. Getting right the color, size, distance, trajectory of something to eat under a variety of contextual conditions is critical to the survival and/or ability to reproduce of any animal with a perceptual system. Representing such matters veridically is a function of the perceptual system.

One can evaluate even a frog or insect's perceptual representation for how accurately it corresponds to the objects, properties, and relations that it represents. This is an evaluation of veridicality, specifically correctness, not an evaluation of truth. Truth is propositional.

Representational correctness is a representational good for a perceptual system. As always, goods imply standards for achieving them. So norms governing the production of representations relative to achieving this good are necessarily involved. The norms that interest me govern production of correct perceptions in normal circumstances, within the perceptual system's perspectival limits. The animal has not mastered these standards. They are implied by the animal's representational abilities set against the representational goods of its perceptual systems.

At this phylogenetic level, I resist counting operating reliably and well with respect to standards for achieving representational correctness an *epistemic* good. I resist counting the associated standards *epistemic* norms. Epistemology concerns knowledge and belief. Knowledge and belief are missing from this area of the animal kingdom. The norms are precursors.

I have spoken of goods for the perceptual system. What of goods for the animal? The fundamental good for most animals is coping with the environment in a way that facilitates survival and reproduction. Other goods for the

¹¹ Cf. Mazokhin-Porshnjakov, G.A. "Recognition of colored objects by insects" *The Functional Organization of the Compound Eye*, in Bernhard ed. (Oxford, Pergamon Press, 1966); Randolph Menzel, "Spectral Sensitivity and Color Vision in Invertebrates", in *Comparative Physiology and Evolution of Vision in Invertebrates: Invertebrate Receptors*, H. Autrum ed. (Berlin, Springer Verlag, 1979); Christa Neumeyer, "Comparative Aspects of Color Constancy" in *Perceptual Constancy*, Walsh and Kulikowski, eds. (Cambridge, England; Cambridge University Press, 1998); David Ingle, "Perceptual Constancies in Lower Vertebrates" in *ibid.*; David Ingle, "Shape Recognition in Vertebrates" *Handbook of Sensory Physiology* vol. 8 Held, Liebowitz, and Teuber, eds. (Berlin, Springer Verlag, 1978); R. Wehner, "Spatial Vision in Arthropods" in *Comparative Physiology and Evolution of Vision in Invertebrate: Intervertebrate Visual Centers and Behavior*, H. Autrum ed. (Berlin, Springer Verlag, 1981); C.R. Gallistel, "Insect Navigation: Brains as Symbol-Processing Organs" in *Invitation to Cognitive Science IV* (Cambridge, Mass., MIT Press, 1996). For an extensive review, see C.R. Gallistel, *The Organization of Learning*, *op. cit.*

animal depend on goals, and functions of the particular animal. Because the perceptual system's operating well is a good for the animal in coping with its environment, and because the perceptual system's primary good is veridical representation, veridical perception is a good for the animal-type, at least on some occasions. The representational successes help explain the content, which is what it is partly because of its contribution to the pursuits of the animal. This argument has an empirical component—that perceptual systems are good for the animal in coping with its environment. I believe that the empirical component is safe.

Whether there is an analogous argument that primitive analogs of *warrant* for the perceptual system (or perceptual representations) are goods for the animal is more doubtful. Primitive analogs entail a reliable system of representation. Empirically speaking, such analogs are commonly goods for the animal. I am, however, interested in an apriori question. Reliability in normal circumstances, a necessary condition of any analog of an epistemic norm, is a feature of *types* of representation. Veridicality is a feature of tokens. Veridical representations are inevitably, on some occasions, good for the animal-type's endeavors, insofar as perception is a good for them. But is reliable representation inevitably a good for the animal, relative to some function or endeavor of the animal? I am doubtful about an affirmative answer at the level of animals that lack propositional attitudes.

Of course, some aspects of some perceptual systems are unreliable. The representation of danger in rabbits is probably highly unreliable, because of a dominance of false positives. *Danger* may be a representation in an action system rather than one produced by the perceptual systems of lower animals. The same point, however, can be made by reference to what are clearly perceptual representations in the narrowest sense. Spatial perception of viable routes around obstacles is apparently real but unreliable in some slow-moving animals, like lobsters or crabs, for whom collisions are not costly and who are not often threatened by prey.¹²

Is perceptual reliability always a good relative to some function or pursuit of lower animals, as opposed to being merely good relative to the representational function of their perceptual systems? In the animal's actual environment and given certain limitations in its capacities, reliability can have unacceptable costs. Reducing false positives might require slowing down the representation of danger until a better view is possible—with fatal results. The relevant norms are geared to the animal's general capacities and normal environment, the one that its representational content is geared to.

Mere coping with the environment to fulfill biological functions and ends does not seem in general or apriori to require reliably veridical perception. If the *animal did* anything that aimed at veridicality, then reliability would nec-

¹² David Ingle "Constancies in Lower Vertebrates", *op. cit.*

essarily be a good for the animal relative to the aim of that activity. The formation of perceptions is not, however, an activity, goal, commitment of the animal in the relevant sense. It is not obvious that perception must be reliable to serve other natural functions or pursuits of a lower animal. Since reliability is a necessary condition on any good that presages epistemic warrant, it is doubtful that the primitive precursors of epistemic perceptual warrant are, a priori or in general, goods for lower animals.

Reliable veridicality is as a matter of empirical fact often a good for lower animals. Veridicality is necessarily a good for the perceptual system itself. Pre-epistemic analogs of epistemic norms for the fulfillment of this good remain applicable.

Representational norms, at this phylogenetic level, bear on the well-functioning of the perceptual system, not specially on the animal's well-functioning. Most failures to meet the norms will be mere malfunctions in the animal's perceptual apparatus. One can, however, imagine conditions in which certain distractions or informational overload lead to misrepresentations. Both sources of error—malfunction and informational overload—contrast with *brute errors*: errors that are compatible with the animal's perceptual apparatus' operating as well as it can. For example, visual illusion as of a moving physical object could be caused by an unusual confluence of light on the frog's visual receptors. Only failures of the former sort are failures to meet the primitive precursors of epistemic norms.

Perceptual systems are associated with characteristic functions and activities of animals. These activities seem to be guided, even in lower animals, by computations on representations by the animal's subsystems. Given perceptual representations of a honey source, and of landmarks, like the sun, the insect's internal systems compute a representation of a direction that enable it to navigate between the source and the hive. The system adjusts spatial parameters as new information comes in. There are standards, given the animal's abilities and given the aimed-at action, for connecting perceptual representations with activity-guiding representations. So there are norms applying to the transitions between perceptual representations and representations that guide the animal's activity.

Suppose that a frog's perceptual system presents the illusion of a moving object, and the frog's tongue moves to serve the frog's needs, in this case without effect. The frog's visual system has operated well and (let us suppose) reliably, given its inherent limitations and given the stimulation that it received. The error is brute, not a failure to meet the primitive analog of an epistemic norm. If the frog's perceptual system is distorted by brain lesions, it regularly malfunctions. If it is given too many conflicting perceptual signals, the frog might act funny. Disease and informational overload undermine patterns that tend to fulfill the perceptual system's representational function.

Hence they undermine the representational good of the perceptual system itself and its functions in serving activities of the whole animal. Such failures are the phylogenetic predecessors of failures to fulfill epistemic norms.

In evaluating a normally reliable lower-animal perceptual system as having done well with respect to the good of representing veridically, one is not evaluating the animal in any way beyond evaluating its sub-system. The transitions between the animal's perceptions and its activities are also, I conjecture, not acts by the animal. The perceptual system is not evidently connected to any mental acts or mental commitments by the whole animal. In this respect also, we may have a distinction between epistemic norms and their phylogenetic predecessors. Still, the framework for evaluating representation and representational systems at this low phylogenetic level provides an illuminating background, I think, for considering representational activity in more advanced animals.

III.

I will assume that primates and other higher mammals are known to have propositional attitudes—beliefs, conceptualized wants, and intentions—as well as perceptions. Having beliefs requires having a capacity for inference—for truth- and reason-preserving propositional transitions among propositional attitudes, transitions that are attributable, as activity, to the whole animal. Simple logical, inductive, and means-end inferences are present in the mental activity of higher non-human animals.

There is here again the representational good, veridicality. There are norms for yielding this good, given the animals' abilities and limitations. The norms govern not only the perceptual system, but formation of beliefs and inferential transitions. In evaluating cognitive systems with propositional attitudes, it seems appropriate for the first time to speak of epistemic norms. I shall do so.

Lower animals may differ from higher animals in that pursuits of the whole animal, as opposed to its sub-systems, involve only overt practical agency. The issue here is whether mental acts become possible only in the context of propositional abilities. This is an interesting and delicate issue. There are certainly non-propositional mental acts. Imaging a past perception does not, I think, have propositional form. But when we think of it as a mental *act*, it is natural to think of it as a product of an intention, which does have (*de re*) propositional structure. I shall not try to settle the delicate issue.

There is empirical evidence of rotations of images in lower animals. Whether these are attributable to the whole animal and whether they can count as acts seem to me to be questions that invite deeper reflection on the distinction between an animal and its sub-systems and on the notion of pre-propositional, pre-intentional animal acts. In any case, in a wide variety of

lower animals, perception serves agency without recourse to mental acts. Perception itself involves mental—or at any rate, psychological—states and events, but not, I think, mental acts. The whole spider acts in eating, as opposed to digesting, its prey. But there appear to be no intervening mental acts by the spider, as opposed to intervening events in the animal's perceptual and act-representation systems.

In higher animals, the relation between perception and activity is more complex. Mental activity by the whole individual—for example, propositional inference—comes to supplement physical activity in being functionally connected to perception. Mental activity helps guide practical agency. A hierarchy of further activities comes into play as perception is integrated into propositional activity. Then the animal's activity can be evaluated representationally and epistemically, as well as practically.

Having propositional attitudes entails having beliefs, including occurrent thoughts that instantiate beliefs. Having such belief and thought entails that the individual has commitments to the belief's, or thought's, being true. What is meant by "commitment" here? The individual does not merely have a hypothetical or suppositional orientation with respect to the belief, and this orientation is such that the individual undergoes a representational failure if the belief is not true. Moreover, the committal orientation involved in belief has implications for commitment to activity, under appropriate circumstances, by the whole individual or animal. Such activity includes committal occurrent thoughts that instantiate the belief, and drawing inferences from the belief. This mental activity is inevitably involved in potential further practical activity—acting on belief. The commitments need not be understood or reflected upon. Their significance is functional. A system of beliefs necessarily has a function of representing truly. This function, and associated goals and commitments, are attributable to the whole individual. Truth is a representational good relative to the relevant commitments. So individuals that have propositional attitudes necessarily have commitments with representational goods of veridicality, indeed truth. The relevant goods are goods for the whole individual. Standards governing such goods apply to the individual's orientations, acts, and states.

As noted, having beliefs entails engaging in committal inferences. Inferences are acts by the whole individual. Carrying out inferences necessarily, not as a contingent matter, entails having veridicality—in this case, truth—as a good, relative to the function of the pursuit. The mental activity by the whole individual involved in having propositional attitudes and engaging in inferences has the function of obtaining and preserving truth. Fulfilling this function is a representational good for the animal. The animal does not and need not undertake this as a project. It cannot avoid having the function. It is

apriori that the animal has this good and this function by virtue of being a believer.

This good and this function are, of course, not the animal's only ones. They need not be the only functions even of the animal's committal thoughts, attitudes, and inferences. The attitudes and inferences may have practical functions—helping the animal find food, avoid predators, locate shelter. What is a good for the animal relative to the representational function of inference to preserve and arrive at truth, or relative to the commitments involved in thoughts and beliefs, may not be a good for the animal relative to the practical functions of inference or belief formation. Certainly an epistemically good inference can be bad relative to a biological function. For all sub-human animals, and most humans, representational goods are less important than other practical goods. Still, it is apriori that having representational functions and goods is a consequence of having propositional attitudes.

Epistemic norms are apriori necessary concomitants of a system of belief and inference having the representational function of producing true representations. Arriving at truth is a representational good and an epistemic good. The inferences that the higher animal engages in are subject to norms governing rational inference—norms for preserving truth. The committal elements involved in inference and in the formation of belief are subject to epistemic norms governing the formation of true conclusions and true beliefs generally. The animal cannot renounce these norms, even if it can conceptualize them enough to renounce them.

In critical reasoners, these norms become objects of reflection. But to be applicable to mental states of non-critical reasoners, many such norms need not be conceptually accessible to reflection. In young children and higher subhuman animals, they are not conceptualized. Such children and animals *cannot* think about them. They lack the concepts *epistemic, warrant, entitlement, reason, reliable, competence, entails, perception, and perceptual state*. It is doubtful that all mature human adults have all these concepts. Thus for many beings with warranted beliefs, fulfillment of the relevant norms is constitutive of an individual's having epistemic *entitlements*—having epistemic warrants that need not be conceptually accessible to the warranted individual.

In being part of a system of belief, perceptual representation becomes subject to epistemic norms. This is partly because perception becomes subject to "top-down" epistemic interference and influence. It is mainly because relations between the perceptual system and the system of propositional attitudes are matters of epistemic import. Perceptual belief and perceptual knowledge depend for their content and warrant on perceptual states. Norms governing well-functioning by the perceptual system for achieving veridicality and for

-serving the formation of true beliefs are directly relevant to norms governing attaining truth in perceptual belief.

IV.

I now want to discuss some elements of the epistemic structure of perceptual belief. I assume the traditional view that perception of an object or property requires a concurrent causal relation to it. Some have thought that since perceptual representations can be veridical of objects that are not perceived, the perceptual representation can be veridical even though no object of the proper sort causes the perception.¹³ They reason as follows. The visual system can represent a scene as containing a particularly shaped surface with a particular color shade. There might well be a surface with exactly that shape and shade elsewhere in the world, causally unrelated to the perception. The instance of the given visual representation could be caused by some other sort of surface, producing a misperception. Then, according to this reasoning, the visual representation would be veridical, but it would not be a perception. Perception but not veridicality, they say, requires that the causal relation be intact.

This view fails to connect perception with its correctness conditions. It is right to maintain that the general elements in perceptual representation can be *veridical of* entities that are not perceived. But it ignores the singular elements—the purported representation of *particular* objects, instances of properties, or token events—in perception. These singular elements can succeed when some of the perceptual representational types are misapplied. And the veridicality—correctness—of the perceptual application of the perceptual types clearly depends on their being applied to particular instances. Thus the scene is perceived, and the singular element applies to it; but the perceptual representation, which includes singular and general elements, is (mostly) not veridical.

¹³ Cf. Alan Millar, *Reasons and Experience* (Oxford, Clarendon Press, 1991). An alternative view of how causation enters into perceptual representation, which I also reject, maintains that the self-referential *and is the cause of this very perception* is part of every perceptual representational content. This is hyper-intellectualized in attributing the category of a perceptual state to a perceptual system. Perceptual systems are not so self-conscious! For perception to connect to its object, the causal relations must be in place. But the perceptual system does not itself specify these conditions between itself and the world in all its perceptual categories. For an expression of this sort of self-referential view, see John Searle, *Intentionality* (Cambridge, England, Cambridge University Press, 1983). I discuss his argument in "Vision and Intentional Content" in *John Searle and His Critics* (Oxford, Basil Blackwell, 1991). Something like this view was defended a half-century ago in psychology by A. Michotte, *The Perception of Causality*, Miles trans. (London, Methuen, 1963, of 1946 French edition). There are a variety of empirical difficulties with the view, and it has not been accepted in psychology. Cf. Vicki Bruce and Patrick Green, *Visual Perception, Physiology, Psychology, and Ecology* (Hillsdale, New Jersey, Lawrence Erlbaum Associates, 1985), pp. 297ff. On the other hand, I believe that many perceptual frameworks, spatial representations for example, have a *de se* anchor. I shall discuss these matters on another occasion.

The point and function of the perceptual system is to put the perceiver in touch with particular situations, particular objects or events, or particular instances of properties or relations. Hence representations that type its successful, veridical, applications involve singular elements.

Singular elements in perceptual representation are context-dependent. Perceptual representations represent particulars as instantiations of types. A given, specific shape type can suffice for individual object recognition. But clearly, perceptual representation through perceptual types does not purport to represent particulars uniquely in a context-free way. The perceptual system is incapable of being committed to there being only one particular in all the universe that looks a certain way. It cannot generalize in that way. There could be exact look-alikes, even of whole scenes, elsewhere in the universe that are not perceived or represented. The perceptual system represents the objects, properties, scene, before it, not the look-alikes. Its function is not simply to produce accurate representations of some entities or other, but to represent the objects (events, property instances, relation instances) that the individual engages with. To fulfill this function, its representational content must be satisfied by particulars of the represented types, which are in the appropriate causal relation. Thus the content must include context-dependent singular representational elements, analogs of singular demonstratives. They are modes of singular representation that depend for successful representation on the causal relation to perceptual objects in context.

Singular elements in perception, referring to objective particulars, are widely overlooked in the analytic tradition. It is often presupposed that singular reference can emerge only through thought. The presence of singular referential elements in perception is independent of the empirical question of what sorts of entities a perceptual system represents. The presence of singular elements, purporting to represent particulars of some sort (these may include instances of properties either in addition to or instead of objects) is established by apriori reflection on the nature and functions of perceptual systems.

Perceptual systems commonly represent at different levels of abstraction. The sense of abstraction that I have in mind is that a representation R0 is at a lower level of abstraction than a representation R1 along a certain dimension if it applies to a proper subset of the perceptual objects that fall under R1 in that dimension. Visual systems group color shades or shapes under generic perceptual rubrics as well as shades and specific-shapes at the finest grain distinguishable by the perceptual system. Color types like *blue* as well as particular shades of blue, and shape types like *vertical line*, as well as vertical lines with particular shapes, orientations, and lengths, are represented.¹⁴ The

¹⁴ Sutherland, N.S., "Visual Discrimination of Orientation by *Octopus*: Mirror Images" *British Journal of Psychology* 51 (1960), pp. 9-18; and "Outline of a theory of visual pattern recognition in animal and man" *Proceedings of the Royal Society, London, B*,

groupings at various levels of abstraction can be tested empirically by testing which discriminated grouping the animal can make using purely perceptual cues. Such grouping occurs in lower animals as well as higher animals, and can be shown not, in general, to involve propositional attitudes.

Perceptual representations are necessarily "committal". Perceptual systems do not represent in a suppositional, neutral, or hypothetical way. The world is presented as being a certain way. This is not necessarily a commitment by the whole animal. It is not intrinsically associated with representational acts by the whole animal, though it has implications for practical activity by the animal.

I want now to give a preliminary characterization of conceptual representation. Like perceptual representations, concepts are intentional or representational, in the sense that their identity and function involve purporting to be about something. Unlike perceptual representations, concepts help type propositional abilities, including abilities of the whole animal to engage in thoughts and inferences. Concepts are essentially sub-propositional components of *propositional* representational contents.¹⁵ Concepts are standing or context-free elements in thought. They have their representational identity independently of a particular context of application.¹⁶

Some concepts are essentially conceptualizations derived from perceptual representations. A *conceptualization* derived from a *perceptual representation* is a concept that takes on the range of reference of a general perceptual

(1971), pp. 297-317; cf. also David Marr, *Vision* (San Francisco, W. H. Freeman and Company, 1982), *passim*. There are other notions of abstraction that are relevant, but I shall not try to expound them here.

¹⁵ I take "intentional" to be equivalent to "representational". I take "representation" to apply to abstract types, not to tokens. As noted earlier, "representation" is for me equivalent to "representational content", as long as content is understood to have intentionality, or representational properties. Representational contents are, or at least function to be, about something. For present purposes, I am neutral about the exact ontology of contents. Representational contents should not be confused with Russellian propositions, which I believe have no direct psychological relevance. Concepts are certain sub-propositional components of *propositional* intentional contents. Concepts help type both the referential characteristics and the cognitive abilities associated with psychological states and events. Kant characterized concepts as predicates of judgments. Cf. *Critique of Pure Reason* A69/B94. There are functional elements in judgments that are not strictly predicative, and there are individual concepts, like those expressed by the numerals and by "God", that are not predicates. But Kant is right to relate concepts essentially to propositional contexts, and the emphasis on predication as the proto-type is well-taken. Predication is ubiquitous in propositional intentional contents. See notes 16-18 for more on my conception of concepts.

¹⁶ Thus I count demonstrative or indexical expressions, like the first-person pronoun, as expressing a concept. The concept is the constant element in all referential applications of the first-person pronoun. To apply to a particular person, the concept must be associated with a token application on a particular occasion. Representations that mark such acts of token application are not themselves concepts. For more on this view of singular reference in thought, see my "Belief *De Re*" *The Journal of Philosophy* 74 (1977), pp. 338-362; "Descartes and Anti-Individualism: Reply to Normore" *op. cit.*

representation and relies essentially on that representation's mode of presentation.¹⁷ The conceptualization can make use of very detailed perceptual representations (the features of a particular face or jagged shape), as well as more abstract ones (face, rectangle).¹⁸

Even those perceptual concepts that are conceptualizations of perceptual representations are more abstract than the perceptions in the sense that they are less connected to concrete sensation-presentations and can occur independently of any perceptual application. They can be consciously used in conditionals, for example, when no perceptual representation occurs, and when none can be brought to mind.

I have been assuming that perceptual representation is not conceptual representation. The fact that lower animals, which clearly lack propositional attitudes, have perceptual states, in a robust sense, strongly suggests—what is already intuitively plausible, from cases like maps, phonograph recordings, and photographs—that not all representation with veridicality conditions is propositional. Since concepts are essentially elements in propositional contents, these types of representation are non-conceptual. I believe that there are strong specific empirical reasons, as well, to believe that perception is non-conceptual and non-propositional. Perceptual representation is organized in fundamentally different ways from propositional organization, and the two sorts of representational content mark fundamentally different sorts of abilities. I think that perception has a topological-like structure. Genuine predication, which is the key element in propositional representation, embodies a type of generality—and freedom from associated application to particulars—not present in perceptual representations, or the other non-propositional representations that I have mentioned. I think that the percept-concept distinction is applicable in human psychology and human

¹⁷ In the basic case the concept is the conceptual counterpart of a perceptual representational type. Then it in effect uses only the perceptual representation, or at least its residue in memory, in the conceptual system. (The perceptual representation can be at various levels of abstraction: very specific shapes, as well as generic types of shape, are represented by the visual system.) For example, there is a basic concept of objects of car-shape that in ordinary circumstances applies to the same things the concept of car does. The concept of objects of such shape is entirely guided by the perceptual system's grouping of what are in fact cars in accord with characteristic visual appearance of shape. But the relevant object-shape concept that is associated with cars has no implications about the objects being cars, having a transportational function, or the like.

¹⁸ I think that the ordinary concepts of pain and of red could not be what they are if they were not constitutively associated (respectively) with the sensation of pain or a perceptual presentation of red. One could not have the ordinary concept of pain, or that of red, if one had not felt pain or seen, with normal color-sighted vision, red objects. Of course, a blind person could have a concept of red. But the concept is different from the ordinary one because the associated cognitive abilities are fundamentally different: the blind person cannot identify or form beliefs about red on the basis of visual presentation. Concepts are representational markers of cognitive abilities.

epistemology. These are complex issues, for discussion on another occasion. Here I simply presuppose the distinction in discussing perceptual entitlement.

V.

Perceptual representation is fallible and subject to correction. Thus I reject what Sellars criticizes under the label “The *Given*”. On the other hand, contrary to Sellars, I take non-conceptual perceptual representations to play a role in the warrant, the entitlement, for perceptual beliefs. The role is, in most instances, more important than collateral evidence or background belief. Normally, perceptual beliefs do not need epistemically relevant contribution from any other source.¹⁹

I take Sellars’ main contributions to be those of sharpening the question of what could count as epistemic warrant for perceptual belief, and of showing that traditional sense-data conceptions of epistemic warrant are incoherent. I think, however, that his main positive view is mistaken. His main positive view is that the only sort of epistemic warrant for perceptual belief must lie within “the space of reasons”, where reasons are understood to be propositional. The positive view is mistaken because it ignores epistemic entitlement. Sellars and philosophers influenced by him have provided no cogent

¹⁹ Cf. Wilfrid Sellars “Empiricism and the Philosophy of Mind” *Science, Perception, and Reality* (London, Routledge & Kegan Paul, 1963), pp. 154-170. Sellars takes “The Given” to satisfy three conditions. The first is that it be a fact that is non-inferentially known. I do not talk of facts. Perceptual belief is, however, non-inferential knowledge of physical objects, properties, and relations in the environment. These are not what Sellars had in mind when he wrote of a non-inferentially known given. But perhaps the letter of this condition is fulfilled. The second condition is that the relevant knowledge “presupposes” no other general or particular knowledge. “Presupposes” is inspecific. I hold that having any propositional attitude requires being disposed to make inferences between it and others. Some of these inferential relations incorporate knowledge. On an interpretation that takes “presupposes” to mean “requires the concomitant existence of”, the condition is not met. I believe that this point is in conflict with the notion of The Given that Sellars had in mind—an uninterpreted infallible acquaintance with some sense-data-like fact. On the other hand, I hold that perceptual beliefs commonly do not rest for their epistemic warrant on any other beliefs. On an interpretation that takes “presupposes” to mean “draws its warrant from other beliefs”, this condition on The Given is fulfilled. The third condition is that the relevant non-inferential knowledge is the ultimate court of appeals for all knowledge, or all empirical knowledge, of the world. This condition is not satisfied. I think that all perceptual belief is defeasible. Sellars’ attack on The Given, like the positions of many who have been influenced by it, is associated with an epistemic internalism about the warrant for perceptual belief that I believe is hyper-intellectualized and untenable. Cf. *ibid.*, pp. 168-70. Sellars’ assimilation of externalism to accounts of thermometer readings overlooks a rich array of intermediate cases, including most psychologically interesting cases of perception and perceptual belief. There are many differences between norms governing psychological states and norms governing thermometers. For example, genuine autonomous perceptual representation is at issue only when the representation plays a role in serving agency in individuals who have the relevant perceptual system. Entitlement is at issue only when the representation plays a role in serving mental agency, including thoughts and inferences.

argument against entitlement as a type of epistemic warrant. They have not considered it.

Philosophers influenced by Sellars have tended to take one of two lines. Some have denied that anything of epistemic relevance occurs in the causal chain between external objects and perceptual beliefs about them. Such views usually appeal to holistic coherence considerations to explain the warrant for perceptual beliefs. By contrast, some have held that perception is itself propositional and provides propositional reasons, in perception, for perceptual belief.²⁰

I regard both of these views as very strange. The one maintains that perceptual states play no role at all in the epistemic warrant for perceptual belief. They are merely causal enabling conditions for there being perceptual beliefs. This view hardly does justice to the epistemic role of perception, and invokes a notion of coherence that has little evident connection with truth, let alone with the role of perception and perceptual belief in ordinary epistemic practice. Perception has a place in the account of warrant because perceptions are representational states whose role in belief formation is subject to normative, epistemically relevant valuation. The other view holds that perceptions are reasons and implies that the formation of perceptual belief is a piece of reasoning—a transition from a reason to what it is a reason for. This view maintains an empirically unsupported and implausible view about the form of perception (as propositional), and its treatment of the transition between perception and belief as a form of reasoning or reason transmission has little basis in intuitive epistemic practice. The upshot of Sellars' position seems to be that either perception is propositional and reason-giving or it plays no distinctively epistemic role at all. What considerations could back such an unattractive disjunction?

Often the view is supported by making alternatives seem pernicious. The alternatives commonly discussed seem to me weak positions, only contingently associated with the point at issue. Sometimes the alternative is seen as the postulation of infallibly known sensations as the epistemic basis for perceptual beliefs. Sometimes the alternative is seen as an appeal to non-intentional sensations as reasons for perceptual beliefs.²¹ I have no sympathy for

²⁰ Philosophers influenced by Sellars who take the first line are Richard Rorty, *Philosophy and the Mirror of Nature*, (Princeton, Princeton University Press, 1979); Lawrence Bonjour, *The Structure of Empirical Knowledge* (Cambridge, Mass., Harvard University Press, 1985); Donald Davidson, "A Coherence Theory of Truth and Knowledge" in *Truth and Interpretation* Lepore ed. (Oxford, Basil Blackwell, 1986). A philosopher influenced by Sellars who takes the second line is John McDowell, *Mind and World* (Cambridge, MA, Harvard University Press, 1994). There are important differences among the positive views of these Sellarsians. Rorty, Bonjour, and McDowell maintain epistemically internalist positions. I am less sure of Davidson on this point.

²¹ These two views mark the opposition as Sellars saw it. At the time Sellars wrote, they were not straw men.

these views. There is no more basic knowledge, much less infallible knowledge, that underlies and supports ordinary perceptual belief. Sensations in abstraction from their representational aspects *could* not be a ground for beliefs about the objective world—as Sellars correctly maintained. *Knowledge* of sensations as such is a late development in child cognition, and certainly not a foundation for perceptual knowledge either.

Some Sellarsians oppose any appeal to a non-propositional perceptual state in the epistemic account of warrant for perceptual belief. They hold that any such appeal is as incoherent as the traditional pure-sensation account, because it grounds epistemic warrant in something that is not a *reason* for belief. No cogent argument for this charge of incoherence has been advanced. No good argument been given for thinking that epistemology must confine itself to reasons and justification.

I certainly agree with Sellars that reasons must be propositional. Perceptual beliefs are not normally reason based. The normative transition from perception to belief is not a piece of reasoning. If perceptual representations were reasons for perceptual belief, such transitions should count as reasoning. But they do not. Such transitions are not justifications in the traditional sense. They are normative elements in entitlement to perceptual belief. Reasons are the products of higher cognitive capacities—capacities for propositional attitudes. The perceptual system does not produce reasons. It does produce representational states that entitle individuals to perceptual beliefs.

The notions of justification, reason, and deliberative reason mark specially important types of warrant. It is nearly trivial that non-conceptual perceptual representations do not provide *those sorts* of warrant for perceptual belief. But these epistemologically “internal” notions cannot do the work of a full epistemology. All epistemic warrants connect to rational inference, since propositional attitudes are constitutively associated with inferences. But reasons are not the only sort of epistemic warrant.

The claim that reason or justification is the only sort of epistemic warrant can be seen as a stipulative restriction on what “warrant” is to mean. But if it is a substantive claim, it hyper-intellectualizes epistemology. It focuses entirely on a kind of epistemic good that derives from the more intellectual aspects of the more intellectual representational systems. In so doing it deprives epistemology of resources to account for more primitive, but nearly ubiquitous epistemic goods. Children and higher non-human animals do not have *reasons* for their perceptual beliefs. They lack concepts like *reliable*, *normal condition*, *perceptual state*, *individuation*, *defeating condition*, that are necessary for having such reasons. Yet they have perceptual beliefs. There is no sound basis for denying that epistemology can evaluate these beliefs with respect to norms governing their formation, given the perspectival limitations and environmental conditions of the believer. There is no sound basis

for denying that epistemology can evaluate their perceptual beliefs for epistemic warrant. There are legitimate questions about animals' and young children's entitlement to their perceptual beliefs.

Even for mature human adults, perceptual beliefs seem normally not to be warranted by reasons. Of course, in some contexts, one can use a back-up reason to supplement an entitlement, or to counter a doubt: "Well, I certainly seemed to see (or I had a perception as of) a such and such". Here the reason is a belief, not a perception. It is a conceptual mistake to think that the perception that is referred to is itself the reason. Reasons are propositional. Moreover, such meta-reasons or retreat-reasons are not fundamental to perceptual belief or perceptual knowledge. An ability to think about appearance or about perception is not necessary for being entitled to perceptual beliefs.

Similarly, in certain contexts, one can justify a perceptual belief by claiming *I saw such and such*. Such a claim can count as a reason. But it can be asked wherein the warrant for such a claim lies. The claim refers to a perceptual state, or presupposes one. Wherein is one warranted in holding this to be a seeing? How is its being a perceptual state (even in cases where it is not veridical) capable of warranting the perceptual belief (which may be mistaken)? It is unacceptable to think that being warranted in one's perceptual belief requires an ability to follow out such a line of questioning. More fundamentally, to be warranted in a visual perceptual belief, one need not have beliefs about one's visual capacities. Such beliefs come later than visual belief about ordinary environmental matters. There is no reason to think that one cannot be warranted in beliefs about the physical environment unless one has meta-beliefs about what capacities one has to perceive it.

Again, such a claim does not indicate the epistemic role of being in a perceptual state in undergirding the perceptual belief. Understanding the warranting force of being in a relevant-seeming perceptual state and having some grip on the general shape of conditions that undermine such force are part of having a reason. The idea that all human adults with warranted perceptual beliefs have even a conceptualized "know-how" mastery of this territory would, I think, overestimate the conceptual capacities of adult human-kind. The idea that adults who cannot conceptualize the relevant epistemic conditions lack warranted perceptual beliefs would, I think, betray a hyper-intellectualized conception of warrant. Such adults lack reasons and justification, but are warranted in their beliefs—entitled to them. Epistemology must acknowledge elements of warrant that are not conceptually accessible as reasons to the warranted individual if it is to give a tenable account of perceptual belief.

I have not discussed all arguments for the view that perceptual states are either conceptual (propositional) or lack epistemic relevance. There have been claims that any opposing view must fail to capture the direct singularity of thought, or must "lead to" scepticism or a "veil of ideas", or cannot distin-

guish human beings from thermometers. I think that many such claims are vague or metaphorical. Many rest on mistaken assumptions about the form of the opposition. None provides good ground to hold such an extreme and unattractive disjunction.

Perceptual beliefs that are true are commonly associated with veridical perceptual states. A claim that such perceptual states are in general propositional is empirically implausible.²² Moreover, there is no plausible basis for thinking that perceptual states, even if they were propositional, provide *reasons* for perceptual beliefs. On the other hand, a claim that the representational perceptual states play no epistemic role in the warrant for perceptual belief loses touch with what is distinctive of empirical warrant—its reliance on sense-perception. Rather than pursue these uninviting lines, I would like to develop further the role of entitlement in the epistemology of perception and perceptual belief.

VI.

I see entitlements to perceptual *beliefs* as fulfillments of normative conditions on perceptual states and on transitions from them to perceptual beliefs. In this section I will develop some of the motivations for my conception of entitlement, and I will discuss the positive contribution of perceptual states to our entitlement to perceptual beliefs. In subsequent sections, I will discuss ways in which the positive contribution of perception to entitlement can be outweighed and the role in entitlement of transitions between perception and belief.

As I noted earlier, norms governing the attainment of veridicality in perceptual systems of lower animals are not epistemic norms. Epistemic entitlements emerge only where perceptual representations are part of a psychological system that involves beliefs and mental acts. In individuals where perceptual systems serve the formation of perceptual belief, fulfillment of norms governing perceptual correctness is a significant element in the cognitive system's fulfilling epistemic norms that govern the formation of true belief. Then fulfillment of such norms governing perception contributes to an *epistemic* good, and can help yield epistemic warrant, or entitlement.

In my view, epistemic entitlements do not attach to perceptual states, perceptual transitions, or perceptual competencies, even in beings that do form

²² There are extremely delicate conceptual issues associated with this claim. There are few, but nevertheless a few, psychologists as well as philosophers who maintain that all representation is propositional. I believe that this position has never been intuitively plausible, and that it has become less empirically reasonable, and more widely rejected in psychology, over the last two decades. For now, I simply assume that it is plausible that perceptual representation is not fundamentally propositional (hence not conceptual). I will try to show the coherence and shape of a larger view that includes this position. I intend to discuss this issue on future occasions.

perceptual beliefs. Perceptual entitlements attach to perceptual beliefs. An epistemic entitlement to rely on a perceptual state or a perceptual system just *is* an entitlement to hold appropriately associated perceptual beliefs. Still, one can reflect on the contribution that perceptual states make to an entitlement to perceptual beliefs.

Not all norms governing an animal's response to perceptual states are epistemic norms. By relying on perceptual competence as practically beneficial, the animal fulfills practical norms. Here too the competence is constitutively individuated and explained partly in terms of its role in guiding practical activity successfully, so there are presumably practical entitlements that parallel epistemic entitlements. Fulfilling *epistemic* norms requires that the reliance on perceptual competence, or on a particular type of perceptual state, further true belief.

What is the contribution of perceptual states *per se* to entitlements to perceptual beliefs? And how are entitlements undermined if this source of contribution falls short?

There are two elements to the distinctive contribution of perceptual states to epistemic entitlements. The fundamental one derives from the anti-individualist account of the individuation and nature of perceptual states. A condition on particular perceptual representational states' having the content that they have is that there have been both causal-formative interactions (which are not in themselves representational) and representationally successful interactions between instances of types of relevant perceptual referents and aspects of the individual's perceptual system (in either the individual's history, or in evolution of the system in his evolutionary ancestors, or in some other way).²³

The other element in the contribution of perceptual states to epistemic entitlements is a condition on the successful representational relations that

²³ "Relevant" must be understood to allow that the individual might have representations for types with whose instances neither the individual nor any of the individual's evolutionary ancestors has ever interacted. For example, one can erroneously perceive an instance of a color shade or shape type as being present, even though no individual has ever interacted with such an instance. This is possible as long as an ability to represent a spectrum of color shades or shapes is in place. Instances of other types in the spectrum must have been interacted with. The relevant ability to mistake a never-instantiated type can be in place as a result of ancestors' interacting with instances of some shades in the spectrum or with instances within a combinatorial array of shape types. The animal is able to respond in a systematic way to differences within the system of representations for the relevant parameter. What is important is perceptual states are individuated in such a way that failures of reference be explained against a background of successes. The routes that such explanation take can be quite complex and indirect. Perceptual states that never were veridical will fail the reliability condition on entitlement to perceptual beliefs relevantly associated with those states. I state the reliability condition in the next paragraph. Perhaps the reliability condition should be relaxed to apply to perceptual states that are reliable in most instances within a spectrum or array.

play a role in the individuation conditions. To contribute to epistemic entitlement a perceptual state (type) must be *reliably veridical in the perceptual system's normal environment*. The normal environment is the one by reference to which the perceptual content of the perceptual state is explained and established. Reliability is necessary for epistemic entitlement because all epistemic warrant is fundamentally an epistemic good inasmuch as warrant is a good route to truth and knowledge. When a perceptual state is reliable in the normal environment, and certain other conditions on entitlement are met, the individual is entitled in *any* environment to perceptual beliefs that are appropriately derived from the perceptual state.

The idea behind this view of entitlement is that reliance on perceptual states is warranted partly because the very identity of the states is constitutively and explanatorily associated with veridical representation—and with guiding activity with respect to entities veridically represented. The role of a state in responding to such entities and in functioning to guide practical activity with respect to them helps make it what it is. Veridicality enters into the very nature of perceptual states and abilities. So it is built into the *nature* of the competence associated with the formation of a reliable perceptual state that the state make a non-accidental, explanatory, positive contribution to true belief and knowledge in the animal's normal environment—in animals that are capable of true belief and knowledge. Being in a reliable perceptual state in itself makes a positive contribution to a believer's having a *prima facie* entitlement to form an appropriately conceptualized perceptual belief from it.

Because a perceptual competence or state is individuated and correctly constitutively explained partly in terms of representational success, exercises of the competence, or instances of the state, have a structural and explanatory connection to veridicality. This connection is in itself an epistemically relevant good in that it contributes to the representational good, truth, that epistemic goods serve. The competence or state's being *reliably* veridical in normal conditions gives this epistemically-relevant good a status sufficient that it can contribute to the specific epistemic good, warrant.

All perceptual states that represent entities in the environment are necessarily individuated through relations to the environment. Individuation of the states and of their representational contents is simultaneously individuation of a perceptual competence associated with those states and contents. Let us suppose the competence reliably veridical. If the reliability of a perceptual representation is not grounded in the individuation and nature of the state, then the reliability cannot yield entitlement. The reliability could be inductively learned by the individual. But then the inductive connection would be the source of warrant—in fact, justification. It would rely for its justification on perceptual beliefs (in the induction base) that drew warrant from the indi-

viduation conditions of underlying perceptual states whose reliability *is* grounded in the conditions under which they are individuated. If the reliability of a perceptual state is not inductively learned by the individual, and is not grounded in the individuation and nature of the state, then reliably veridical perceptual states make no contribution to entitlement. The individual would have no entitlement to rely on the *de facto* reliable state or competence.

There is a difference between the individuation condition and the reliability condition in the way they bear on epistemic entitlement. Only the individuation condition determines the nature of the individual's competence, and helps determine norms governing exercising the competence well. The reliability condition bears only on what sort of norm is being fulfilled. It bears on whether the standards that the optimal exercise of the relevant competence meets are standards for yielding warrant or epistemic entitlement. The individual might meet certain epistemically relevant standards—representing as well as possible given its limitations, for example—which are not sufficiently demanding to be standards on epistemic warrant. Some aspects of the individual's perceptual apparatus may be constitutionally insufficiently reliable to contribute to warranted belief, but epistemic norms less demanding than those governing warrant would still be applicable.

The role of reliability in the account is relatively straightforward. It underwrites the fundamental function of epistemic entitlement—serving as a good route to truth. The role of the normal environment in determining entitlement is more complex and interesting. The *normal environment* for a perceptual state or competence is that environment, or those aspects of an environment, by reference to which the nature of the perceptual state or competence is explained. Environments that are intuitively different but which are relevantly similar, in ontological and epistemic respects, to the content-establishing environment count as the same environment for our purposes.

One can ask why should reliable success in such a normal environment bear on the contribution of perceptual states to entitlement in abnormal environments that the individual may have landed in. From a certain Olympian point of view, one can wonder why the normal environment should have a privileged role in determining warrant—in determining what it is to be a good route to truth—when it is the individual at a given time for which warrant is supposed to provide a route to truth. The individual might, from the Olympian point of view, be in any number of possible “abnormal” environments. In many such environments, a perceptual state or competence, which, on the present account, contributes to entitlement, would not be a good guide to veridicality or truth. It would not contribute to a tendency to form true perceptual beliefs in that environment. Why does the normal environment have privileged status? Why does its contribution to entitlement carry over to other

environments that the individual might form perceptual beliefs within? Why are the individual's perceptual entitlements backed, in effect, by some bias or presumption—not necessarily conceptualized by the individual—that he is in the normal environment, in the absence of grounds to think otherwise?

The questions are about entailments of the conjunction of the two conditions, the reliability condition and the individuation condition. I want to explain the consequences of these two conditions by starting from more fundamental considerations. In effect, I want to give a transcendental argument that shows why these consequences hold. The argument explains in a deeper way how normal conditions enter into epistemic warrant.

The short answer to the question why the normal environment is privileged is that the reliability of the perceptual state or competence in any other environment would be accidental relative to the nature of the state or competence. Reliable connections to the world that are accidental relative to the conditions that individuate the individual's perceptual states and competencies contribute nothing to empirical epistemic entitlement. Something in the nature of the individual's perspective, ultimately including perceptual beliefs, must tend toward true belief in conditions that individuate the nature and function of the perspective. I lay out an argument for this standpoint over the next nine paragraphs.

A condition on warrant is that it be a good route to truth. The truth of perceptual beliefs is grounded in the veridicality of the underlying perceptions. So to contribute to the warrant of a perceptual belief, a perceptual state must be a good route to veridicality. Warrant, hence entitlement, attaches to perceptual beliefs when those beliefs derive, appropriately, from perceptual states that contribute in a positive, reliable way to the formation of true belief.

A further condition on warrant is that the connection between warrant, including entitlement, and veridicality be principled or explanatory. Epistemic entitlements are goods that partly depend, in a principled or explanatory way, on the nature of the individual's representational competence or perspective—hence on the nature of the individual's perceptual states.

Any contribution that a perceptual state makes to perceptual entitlement must depend on its having a *nature* that is reliably veridical. Although perceptual states must be systematically connected to other perceptual states to be what they are, and although the positive contribution of a perceptual state to entitlement can be undermined or overridden, a perceptual state makes a contribution to entitlement that is not derived from positive contributions of other states. An empirical contribution derives from the state's having a nature, as marked by its representational content, that is itself reliably veridical. For reliable perceptual states in individuals with propositional attitudes, a defeasible positive contribution to epistemic entitlement is coded in the per-

ceptual state-type (given that it is relevantly reliable). Having a reliable perceptual competence with respect to a particular perceptual state makes that state, through its nature, part of a good route to truth. This point reflects the fact that the occurrence of a percept, functioning normally within a cognitive system, *in itself* makes a defeasible contribution toward one's being entitled to a belief that is appropriately related to it.

So for perceptual states and competencies, the explanation of the relation between warrant and veridicality must hinge on a principled connection between the natures of the perceptual states, on one hand, and their success in representing the world, on the other. The explanation cannot fix on some feature of the world that is accidental relative to the natures of the states. In other, non-normal conditions the reliability of reliable states cannot be explained in terms of relations to the environment that are coded into the natures of the states themselves. Reliable connections to the world that are accidental relative to the conditions that individuate the individual's perceptual states make no contribution to empirical epistemic entitlement.

This fundamental feature of warrant is underwritten by the connection between perceptual entitlement and the anti-individualist natures of perceptual states. That connection explains the reliable success of those perceptual states that *are* reliable, in terms of the environmental conditions and environment-state relations that determine the natures of those states. Reliability in *those* conditions is principled and explanatory relative to those natures. It is connected to the conditions under which the natures of the perceptual states are determined. A corollary of this point is that reliable connections that enter into the individuation of perceptual states are not happenstance. They are law-like connections between psychological state and the world that have the further feature of helping to constitute or individuate the perceptual states by determining what representational content they have.

Because of every individual's perspectival limitations, and because of the nature of empirical warrant, every perceptual state, including those that contribute to warrant (and likewise every perceptual belief, including warranted ones) could in principle have been in circumstances in which it was prone to error. No perceptual state type can be reliably veridical under all conditions. Perceptual states are components of perceptual competencies. Our perceptual competencies are fallible. All perceptual states (type or token) and all perceptual competencies are subject to possible error. In principle, any perceiver could be placed in a situation in which any given perceptual state-type could be regularly mistaken—if the perceiver already has perceptual states with definite content. In other words, every perceptual state type would be unreliable in a variety of possible abnormal environments that are perceptually indiscernible, at the time of a given perception, from the normal environment.

It is a fundamental feature of perceptual *warrant*, hence perceptual entitlement, that it allows that an individual can be fooled while retaining warrant. Warrant is retained if the error is accidental, relative to the nature of the perspective. Perceptual errors or unreliabilities that are perceptually indiscernible and derive from brute abnormalities in the environmental conditions that cause perceptual states do not undermine warrant. Both unreliability and reliability in conditions other than those that played a role in explaining the nature of the perceptual state and the exercise of the perceptual competence are accidental relative to those natures. So reliability and unreliability in *such* conditions are irrelevant to the connection between warrant and veridicality. The *only* reliability that is relevant to the contribution of perceptual states to perceptual warrant is one that attaches to the states' normal functioning in the conditions that explain their natures. That is the only reliability that is non-accidental relative to the natures of the perceptual states.

Thus the normal environment—the conditions in which content is explained and established—is privileged in explanation of entitlement. Its privilege derives from the fact that it plays a central role in making the individual's states what they are—a role that abnormal environments do not play. This privilege entails the irrelevance to entitlement of reliability and unreliability in other conditions. So it extends to indiscernible abnormal environments that the individual might contingently find himself in.

The account of perceptual entitlement *must* take as privileged the normal environmental conditions that help individuate perceptual states and competencies. For they alone connect reliable representational success to the natures of the states in a principled way. Reliability or unreliability in any other conditions is accidental relative to the nature of the states. And, as noted, warrant requires of warranted perceptual states and competencies that their being part of a good route to truth be explained in terms of their natures. Although the natures of the states depend not only on relations to the environment, but also on relations to other representational states and competencies, these latter relations cannot by themselves explain representational success, much less reliable representational success.

The presumption of normality in the absence of defeating conditions or defeating considerations is the hallmark of defeasible empirical warrant. Relevant alternatives are considerations that legitimately threaten the presumption of normality, local or global. Thus the Olympian point of view must be tempered by consideration of the role of kind-determining starting points in determining epistemic warrant.

I believe that there are further issues about this presumption or bias. I believe, for example, that the presumption is a condition on the possibility of the successful exercise of any representational competence. Such a condition bears on the nature of entitlement and warrant. But the objective deduc-

tion of perceptual entitlement that I have just outlined will suffice for the present.

The foregoing explanation of the intuitive basis for perceptual entitlement is not meant in itself to speak to scepticism. It does not purport to provide the most global explanation and justification of perceptual entitlement. But it does provide, within the context of a normal scientific and common sense view of the world, a partial but principled explanation of why the perceptual states that in fact *do* contribute to our entitlement to perceptual beliefs count as contributing to perceptual entitlement.

VII.

I move now from the positive contribution by perceptual states to ways entitlement can be undermined through their not contributing in the right way. Insofar as a perceptual state is not induced purely through the exercise of a perceptual competence, but is intruded upon by other representational states, the perceptual state may fail to contribute to epistemic entitlement. For example, if a perceptual state is partly the product of an expectation that distorts normal formation of the state, the state may fail to contribute to entitlement. This failure may occur no matter how reliable the perceptual state is when it is the product of the normal exercise of perceptual competence.

A more interesting failure of a perceptual state to contribute to entitlement is through failure of the reliability condition. The rabbit-danger case from section II is a paradigm case. The perceptual system or the perceptual state type is unreliable in the relevantly normal environment.²⁴ The reliability con-

²⁴ There is a subtle question here about malfunction. The question is whether the entitlement lapses if the system malfunctions in the particular circumstances, but the perceptual system and the particular perceptual state in the individual remain generally reliable. Suppose that, on one occasion, in the normal environment, a ray of the sun, included in the stimulation of the retina, causes a momentary neural malfunction that produces a perceptual state inappropriate to the environmental circumstances, and that the individual could not discern this case from a normal case of having the perceptual state-type. (I am assuming that genuinely hallucinatory states—caused internally without going through the usual visual pathway—are not perceptual states at all, however indiscernible they may be to the individual with the hallucination.) Suppose that the malfunction does not undermine reliability. This is not a brute error, because the fault lies in the individual's subsystem. But the fault is a malfunction, not a misuse. It is certain that the perceptual belief deriving from the perceptual state cannot constitute knowledge. Even if there were compensating malfunctions that happened to produce the perceptual state in the particular case in such a way that the perception was type-veridical, the individual would not have knowledge. The deviant causal chain and malfunction would undermine knowledge.

Would the individual's epistemic entitlement lapse? Internalist intuitions support the view that it does not lapse, on the ground that in responding to the perceptual state in an affirmative way, and in lacking any contrary reasons, the individual has done the best he can on the basis of the perceptual representations available to him. I do not want to rely upon such intuitions. I think the entitlement in ordinary cases does not depend on how the individual does, but on the natures of his subsystems and states, and on the connection of these natures to the objects of perception. In the present case, we have a momentary

dition would also fail if a perceptual state never represents veridically but is individuated through its systematic connection to others that do. Cf. notes 7 and 23.

Barn-type cases, where an individual perceives correctly, but is lucky in doing so because of locally abnormal conditions, undermine knowledge but not necessarily entitlement.²⁵ Similarly, if one is transported to an abnormal twin or vat environment, with no warrant for suspecting this, one can be entitled to rely on one's perception, even if it has become contextually unreliable. This would be because one's perceptual representations are reliable in their normal environment.²⁶

malfunction that (by hypothesis) does not affect the reliability of the system or the nature or reliability of the perceptual state-type. On the one hand, the particular path that normally leads to the perceptual state has been circumvented. This suggests that the entitlement might be undermined. For the functioning that supports reliable belief is not present. On the other hand, the perceptual state-type is what the entitlement derives from. It is produced by a reliable perceptual system; it is itself reliable in normal conditions, and has a nature that normally connects it to appropriate perceptual objects. This suggest that the entitlement remains intact, in something like the way it does in the case of brute error. An unlucky token event blocks knowledge but not warrant.

I am inclined to support this latter intuition, though the notion of warrant may be vague with respect to such cases. I believe that warrant derives from representational states and processes on the basis of their *types* in a given individual's cognitive system—assuming an absence of *representational* states that through their use or misuse affect the warrant of token instances of those types. Since the perceptual state-type is by its nature associated with veridical representation and since it is reliable, the individual is entitled to rely upon it. Token events that are momentary malfunctions in a reliable perceptual system and that produce a reliable perceptual state undermine knowledge but (I conjecture) not entitlement. I believe that momentary malfunctions that cause beliefs that are indiscernible from perceptual beliefs but which do not derive from a perceptual state, produce hallucinations which are not perceptual beliefs. I am inclined to say that the individual lacks any entitlement to rely on such hallucinations inasmuch as they are not products of a reliable cognitive system—in particular, not products of the perceptual system—and are not connected to such products by any rational means. All of these issues, however, deserve more reflection. I owe the impetus for these points to an objection by Tony Anderson.

²⁵ Alvin Goldman, "Discrimination and Perceptual Knowledge" *The Journal of Philosophy* 73 (1976), pp. 771-791.

²⁶ The distinction between successful functioning and reliable functioning must be stressed. Some background of successful functioning is necessary to counting as a perceptual system. The functioning of the rabbit's predator-spotting system is both practically and, in crucial cases, representationally successful in the environment in which it evolved to save rabbits from hawks. But the system is never reliable in spotting hawks. (I think that a perceptual system cannot have modal categories like *possibly a predator*.) Although explanation of successful functioning is commonly in terms of design, in actual fact evolutionary design, I think that an accidentally produced "swampman" might come to have representational content about its environment through interacting with it. I do not mean the notions of function and design to be entirely inflexible or reductively basic. For interesting discussion of these issues, see Alvin Plantinga, *Warrant and Proper Function* (New York, Oxford University Press, 1993), *op. cit.*, and Peter Klein, "Warrant, Proper Function, Reliabilism, and Defeasibility" in *Warrant in Contemporary Epistemology*, Kvanvig ed. (Lanham, Maryland; Rowman & Littlefield, 1996).

There may be more natural cases than the sceptical scenarios in which an individual lands in a globally abnormal environment. An animal's perceptual system may have originally evolved in an environment where certain perceptual representations were reliable. The formative environment may have changed to the point where perceptual representations deriving from approximately the same proximal stimulations are no longer reliable for the individual animal. We can, of course, imagine cases in which the individual's representational content changes to adjust to causal relations with the new environment. But we can also imagine that the practical dispositions of the animal are so geared to the perceptual system (whose basic contents are, we may suppose, evolutionarily determined) that the representational content remains the same. The perceptual system could have evolved to reliably indicate a fine-grained food of a certain sort. The environment may have changed so that fine-grained dust that is perceptually indiscernible to the animal more frequently causes the stimulations than the food. The animal species continues to ingest the food and dust, but the dust predominates, even though it is perceived and treated as food. Our individual animal might survive because the dust is not harmful; the animal gets occasional nourishment from the less common, indiscernible food; and there are other food sources for the animal. Alternatively, the animal and its species might quickly vanish.

Consider the period after the environmental change when the animal relies upon perceptual representations that are unreliable in its current environment. Do the perceptual representations contribute to an entitlement to relevant perceptual beliefs, even though they are unreliable in the animal's present environment, because the representation types in the relevant perceptual-belief system were reliable in the system's original environment? Is the animal epistemically entitled to rely upon those perceptions as veridical?

There are a variety of cases here, including borderline cases. I think that most such cases should be seen on an analogy with sceptical, demon scenarios, in which the individual lands in an indiscernible, abnormal environment. Although the new environment in which the particular perceptual beliefs are unreliable may, in a sense, be the animal's "normal" environment, it is, by hypothesis, not the environment relative to which the perceptual content is explained. It is not the environment to which the animal's practical activity, through its evolutionary history, was ingrained. It is not the environment to which that activity bears a principled explanatory relation. Where a perceptual system yields perceptions and beliefs that are or would be reliably veridical relative to that privileged environment, it contributes to defeasible warrant, in any environment—unless defeating reasons are available.

Of course, no such animal would have knowledge in the relevant cases, even if its warranted perceptual beliefs are true. Even if the animal had a true and warranted belief that it was approaching food (supposing dust is more

common), the belief would not constitute knowledge. This sort of Gettier case is not a matter of the animal's relying on false lemmas. It is a matter of the belief's being "accidentally" true relative to the particular situation in which the animal is embedded.

VIII.

I turn from the role of perceptual states in entitlement to the role of the transition from those states to the formation of perceptual belief. The transition is not necessarily temporal. The perceptual state and the perceptual belief are different psychological states, even if they are caused simultaneously. The transition is certainly not normally conscious or active. It is normative. That there is a transition, and norms governing it, is evinced by the possibility of error in the transition, and by the fact that warrant can lapse if the transition is not negotiated well.

The transition's contribution to entitlement is through its preservation of the contribution of perceptual states to entitlement. To support entitlement to perceptual belief, the transition must reliably preserve the contribution of reliable perceptual representations, and it must do so in a way that leaves the perceptual belief reliable in normal environmental conditions. Moreover, in particular cases, if the perceptual states are veridical, the transition must lead to true perceptual beliefs.

To preserve the contribution of perceptual states, the transition must meet three further conditions. First, the perceptual belief that results must be caused or causally sustained by the process that yields the perceptual representation and the conceptualization derived from that representation. Second, standard competencies must be exercised in the conceptualization and the formation of the belief. These competencies might be inaccessible to reflection, and subject only to limited and learned control. The causal relation must not, however, be the result of fluke malfunction, or manipulation from another being. Third, the transition must fulfill certain further norms that I shall soon come to.²⁷

²⁷ A similar transition from non-conceptual perceptual representations to conceptualized beliefs is postulated in the work of Christopher Peacocke, *A Study of Concepts* (Cambridge, Mass., MIT Press, 1992) and Alan Millar, *Reasons and Experience*, *op. cit.* There are differences. I do not accept the view, held by both, that the non-conceptual perception provides a reason. I do not regard perceptual states as reasons for anything, since they are not propositional. Reasons are the product of a higher faculty and involve a type of generality not present in perceptions. I do not even see the transition as like inference in preserving warrant, since I do not regard perceptual states themselves as warranted or unwarranted. I also think that neither epistemic justification nor entitlement attaches to transitions between such states and perceptual beliefs, although such states and transitions can be aspects of the warrant for, or entitlement to, perceptual beliefs. (Of course, other epistemically relevant goods and norms can attach to these states and transitions.) Epistemic warrant or entitlement attaches to beliefs only. The transitions do show an analog of preservation of truth: the veridicality of a perception can be "preserved" in the truth

There are two important elements in the transition. One is the move from a perceptual representation, which is non-conceptual, to a propositional representation: from presented non-propositional correctness to putative truth. The normative and epistemically relevant character of the transition is not affected by its being a transition from a non-conceptual representation to a propositional one. The most salient aspects of this element in the transition is the association of concepts with perceptual classifications and the association of demonstrative elements in the propositional representation with some of the singular context-dependent elements in the perceptual representation. When the transition goes well, singular reference is preserved in the transition between the singular elements, and reference to properties or relations is preserved in the transition between the general elements. Such preservation of reference lies at the base of the preservation of veridicality. I want to focus here mainly on the transition from general perceptual representations to concepts.

When the transition goes well, the perception is correctly conceptualized. It will be recalled from section IV that a *conceptualization* derived from a general perceptual representation is a concept that takes on the perceptual representation's referent and constitutively relies on its mode of presentation. Conceptualization is part of normal conversion of perceptual representations into propositional representations. It is *conceptual* in that it marks propositional use and propositional abilities. Through conceptualization the simplest

of the corresponding perceptual belief. There are other areas of disagreement. I do not agree with either author in holding that perceptual representation requires some relation to a conceptual system. (On this point, Peacocke has now revised his view.) As against Peacocke, I do not agree that the content of perceptual states is fixed in terms of contribution to knowledge. Many animals have perception but lack propositional attitudes. Hence they lack concepts, and they lack any capacity for knowledge. Moreover, the content of some perceptual states is to be explained partly in terms of relations to the environment that are necessary for practical survival (or survival for reproduction), but which are systematically unreliable—as in the rabbit case discussed earlier. Peacocke's account is more focused on types of correctness conditions than on warrant. Although I do not accept all aspects of this account, I think it a fine contribution. Additional differences with Millar are these. First, Millar seems to think of the initial state as a merely a phenomenal state, not one with representational content and veridicality conditions. I think that this view makes his whole account of the way perception helps warrant perceptual belief difficult to make sense of. Second, I regard Millar's account as hyper-intellectualized in that it requires that if the transition is to yield warrant for perceptual belief, the individual must have a warranted belief that there are no defeating conditions. I will criticize this requirement when we come to the issue of defeating reasons. Third, in the apparent interests of epistemic internalism, Millar rejects the idea, which I take to be indispensable, that epistemic warrants are reliable indicators of veridicality. There is much, however, in the detail of Millar's book that I find insightful. For further criticism of it see Steven Reynolds, "Knowing How to Believe with Justification" *Philosophical Studies* 64 (1991), pp. 273-292. Unfortunately, Reynolds joins Millar in conceiving of perceptual experiences as non-intentional.

sorts of perceptual beliefs are formed—beliefs that make reference to the same objects, properties, and relations that the perceptual system represents.

The perceptual system has a very limited array of general representational types. There is a wider array of perceptual beliefs that do not utilize concepts that are conceptualizations of perceptual representations. We apply concepts like *second-cousin*, *ball*, *amplifier*, *cancer*, *shady character*, through perception even though the perceptual system *per se* does not represent these properties or relations. These concepts are not strictly speaking conceptualizations of perceptions, since they apply to kinds that the perceptual system has no general perceptual representations for. I shall return to these cases. Correct conceptualization not only preserves reference. By constitutively relying on the perception's mode of presentation, the individual's conceptual application is individuated through the same relations to the environment that the conceptualized perceptual state is. Insofar as the perceptual state is reliable, applications of the concept in perceptual belief assimilate the perceptual state's contribution to epistemic entitlement. So in normal cases involving simple conceptualizations, the transition preserves the contribution to entitlement offered by perceptual states.

The other element in the transition is from committal presentation by the perceptual system to commitment by the individual. A notion of commitment by the perceptual system would be somewhat metaphorical. As I indicated in section IV, my term "committal" marks the fact that perceptual systems function so as not to present perceptual representations in a neutral, suppositional way. Perceptual representations count for the animal as having a positive valence. They are presentations as of how things are. The positive valence is shown in the fact that perception normally functions to guide animal activity. At lower phylogenetic levels, this activity is probably mere practical coping with the environment.²⁸ At higher levels, committal presentations from a perceptual system normally engender beliefs, thoughts, and inferences. These are commitments by the whole animal. Cf. section III. Perceptual belief is commonly a default position. So commitment involved in the belief is not necessarily an act, but rather the absence of a withholding of belief.²⁹ But belief is constitutively associated with mental acts and, in the case of perceptual belief, with potential practical activity. In any case, there is

²⁸ Even among lower animals—which lack propositional attitudes—the commitment by the perceptual system might—in special cases—not be matched by a commitment in the whole animal's practical activity. The animal might learn through conditioning not to follow up on a perceptual representation in the way that the animal was initially geared to do.

²⁹ See Daniel T. Gilbert, "How Mental Systems Believe" *American Psychologist* 46 (1991), pp. 107-119; "The Assent of Man: Mental Representation and the Control of Belief" *Handbook of Mental Control*, Wegner and Pennebaker, eds. (Englewood Cliffs, NJ; Prentice-Hall, 1993).

a normative transition from perceptual presentations to whole-individual commitments.

I now consider ways in which the transition's contribution to epistemic entitlement can be undermined. The simplest ways have already been mentioned. The perceptual belief might not be causally connected to a relevant perception. Then it would employ perceptual concepts, but would not be causally sustained by a perception. Or, the perceptual belief might be caused partly by a perception, but might also be manipulated by a further being or be the result of some fluke malfunction. Then the perceptual belief would not be formed through the exercise of a normal competence. In both these cases, entitlement to the resulting perceptual belief would be undermined.³⁰

Entitlement can also be undermined in the transition through haste, bias, or through the individual's misuse of the perceptual representation. It might be undermined through intrusions from elsewhere in the individual's representational system. Insofar as such factors reduce the role of the perceptual state in the formation of the resulting perceptual belief, they tend to undermine entitlement by failing to preserve the would-be contribution of the perceptual state to entitlement to the perceptual belief. Lying alone in the desert, an individual could form the perceptual belief that a con-specific is approaching. A more disinterested, normal, and accurate response to the same perceptual representation would produce a conceptualization of a swirl of yellowish particles. Empirical theory can distinguish between cases where the perceptual representation has itself been distorted and cases where it is left intact, but misused within the individual's broader representational systems.

The individual can have a perceptual representation but withhold a belief in the propositional representation normally associated with it. Animals can

³⁰ This restriction contrasts with my inclination not to count fluke malfunctions as undermining the contribution of perceptual states to epistemic entitlement. (Cf. note 24.) I believe that these matters invite further reflection. My reason for treating the two cases differently is this. The account of warrant (epistemic entitlement) centers on the character of the individual's representational perspective, and on norm-governed transitions between elements in this perspective. The formation of the perceptual state is formation of the perspective. The formation of a perceptual state may be locally abnormal without affecting the epistemic goodness of the perspective. Entitlement depends on the nature of the perceptual states and on norm-governed transitions between them and other representational states. It is true that the malfunction leading to the perception occurs within the individual's psychological sub-systems, and thus contrasts with brute abnormalities in the environment that lead to ordinary brute error. But the malfunction is still "outside" the individual's perspective. It occurs prior to the formation of the first representational states, perceptual states, that help form the individual's perspective. If a relevant, norm-governed transition within the perspective—say, an inference or a transition from perception to belief—malfunctions, then there is a gap or failure in the perspective itself. Such malfunctions bear directly on whether the representational perspective is fulfilling epistemic norms. Obviously, there are a number of epistemically relevant goods and evils at issue. For many purposes it is not important which ones count as bearing on epistemic warrant. I make these points with some diffidence.

learn not to trust their perceptions through a history of bad results in special situations, or even perhaps in ordinary conditions in their normal environments. This element in the transition that passes from a “commitment” by the perceptual system to a commitment, in belief, by the whole animal is a further source of possible failure of entitlement to perceptual belief. Certainly humans, and probably some higher animals, can acquire reasons for not relying on certain sorts of perceptions in certain circumstances. Entitlement is overridden if the individual has a warranted belief—or is neglectful in not having a warranted belief—that the situation is abnormal, or that other defeaters are in play. To have the entitlement, the individual need not, however, have a warranted belief that defeaters are not in play. Such a requirement would make perceptual entitlement unattainable for higher animals and young children. It is enough that the individual lack a reason to avoid relying on the perception.³¹

In summary, entitlement can fail to accrue to a perceptual belief from a perceptual state that is perfectly reliable in the normal environment if there is malfunction in the transition or elsewhere in the cognitive system that affects the normal formation of belief from the perception. It can fail to accrue because through haste, bias, or other some psychological source, the perceptual state is not correctly conceptualized: the normal, reliable competence in forming perceptual belief from perception is not utilized. Entitlement can also fail to accrue because, despite correct conceptualization of the perceptual state, the individual has background reason not to form the belief associated with the conceptualization. Usually such reasons center on local unreliability of the perceptual state, or simply on locally abnormal conditions that affect veridicality.³² When a reliable perceptual state engenders a perceptual belief via normal conceptualizing competence, and no relevant interferences undermine the entitlement, and no background reasons run contrary, then the perceptual state yields entitlement to the perceptual belief.

³¹ The relevant-alternatives epistemology for perceptual belief has been with us for a good while. Cf. Fred Dretske, “The Pragmatic Dimension of Knowledge” *Philosophical Studies* 40 (1981), pp. 363-378. The idea rests on a number of compelling examples. An analysis of what principles govern the examples is still lacking. There is a need for a more general account of why defeaters are defeaters, or more precisely, why relevant alternatives are relevant. I think one should go further than the usual generalized remarks about context-dependence.

³² In principle, though, such reasons could center on the unreliability of the perceptual state in the environment that helps explain its individuation. I believe that it is an empirical question whether all animals with perceptual knowledge, and entitlements to perceptual belief, are capable of reasons capable of blocking the contribution of reliable, well-formed perceptual states to entitlement.

IX.

Before concluding I want to say a bit more about the transition from perception to belief. I have focused on the simplest perceptual beliefs, those that depend on perceptual concepts in the most restrictive sense. These are beliefs that involve *conceptualizations of representations in the basic perceptual system*. I have in mind beliefs involving such concepts as those for shape, position, motion, texture, color—not concepts like *tomato, amplifier, second-cousin, shady-character, sonnet*.

A conceptualization of a perceptual representation is individuated partly through its normal association with the representation it conceptualizes. Both can be more or less abstract in the sense discussed in section IV. The perceptual state is individuated partly in terms of relations to kinds in the world that figure in successful perceptual representation. The concept is individuated through its relation to the percept, and ultimately through relation to the empirical world. A visual concept of a cube is associated with a visual perception as of a cube. The concept and its associated percept are partly individuated by their causal relation to instances of the shape, or to instances of other shapes which are systematically associable with it by combinatorial capacities of the perceptual system. The conceptual content is partly fixed by its marking a propositional ability essentially involved in recognizing the shape through visual representations of it.³³

There is considerable empirical evidence that perceptual systems *per se* make reference to properties within a relatively confined range. The human visual system makes reference to particular colors, shapes, positions, motions, textures—and to certain abstractions from them—as *such*, but not to tomatoes, mercury, cancer, cars, presidents, sonnets, as such. Of course, we are capable of perceptually recognizing tomatoes, cars, and so on. The perceptual system provides perceptions of objects as having certain shapes and colors. The relevant shapes and colors typify tomatoes and cars. The abilities marked by concepts of tomatoes and cars are open to considerations about internal structure and function, respectively. The concepts are also associated with a wide range of other inferential abilities, many of them inductive. Concepts like *fruit* and *vehicle of transportation*, like *tomato* and

³³ As noted earlier, many exact-shape and perhaps some color-shade representations that occur in an individual's perceptual system may never before have been caused by the exact shape or shade, even in the history of the species. But the shape representation will be systematically associated with other shape representations—those of line, surface, curve, edge, volume, and so on—which have been causally associated with instances of the shape represented. So we perceptually represent new shapes by forming the representation from representations that mark perceptual capacities whose nature is individuated partly through causal relations to instances of the shape that they represent, where these causal relations may have occurred in the formation of the perceptual system in the individual's species ancestors, rather than in the individual's history. Similar points apply to representation of color. Cf. notes 7 and 23.

car, which do not conceptualize perceptions, are associated with perceptual recognitional abilities, ultimately grounded in perceptions and concepts of shape, size, color, and so on. The concepts of tomatoes and cars are associated or matched, in learned but automatic ways, with these percepts and with concepts of such shapes and colors. In fact, normally though I suppose not necessarily, learning the concepts of tomato or car includes learning to make these associations. This association appears to occur in high-level vision—the interface between the perceptual system and a wider range of empirical concepts than those that conceptualize purely perceptual representations.³⁴

Although beliefs about objects with certain shapes and colors are not as interesting or salient to our conscious lives as beliefs about objects categorized under ordinary sortals, there is common sense and empirical reason to think that we form such beliefs in the process of forming beliefs about tomatoes, balls, and the like. If something that we had believed to be a tomato were shown to be a fake, we would normally retain a belief that there was an object with the relevant shape and color; and we would recognize that we had held that more restrictive belief all along.

What are we to say about entitlements to perceptual recognitional beliefs that use concepts of entities not specifically referred to as such by representations offered by the perceptual system? Such beliefs seem subject to a wide range of collateral background information. We can be warranted in the perceptual belief that *that* is a tomato if we merely see a green stem peeping out from behind the edge of a box. We can be so warranted if we have background information indicating that the box contains tomatoes. Of course, collateral background information can affect even perceptual judgments involving purely perceptual concepts—involving the concept of a square, for example. One could see no more than a corner of a square but be warranted in believing that that is a square if there were relevant supplementary background information. So collateral information is potentially relevant in both cases.³⁵

So the account of entitlement for perceptual belief, even for more restrictively perceptual beliefs, must be made more complex. There is still this element of simplicity. *In any warranted belief whose warrant comes even partly from present perceptual experience, the contribution to the entitlement that comes from present perceptual experience must go through*

³⁴ Shimon Ullman, *High-Level Vision: Object Recognition and Visual Cognition*, (Cambridge, Mass., MIT Press, 1996).

³⁵ I think that there may be something to the intuition that the less purely perceptual concepts open additional ranges of possible sources of collateral information. For example, background beliefs about functions will be relevant to identifying cars, but not to identifying shapes, whereas background beliefs about shapes will be relevant to both. But I doubt that there is any sharp sense in which the more purely perceptual beliefs are subject to *fewer* (as opposed to fewer types of) complications from background collateral information.

present perceptual states. The entitlement both to perceptual beliefs involving conceptualizations of these perceptions, and to perceptual beliefs involving concepts—like tomato and car—that are not mere conceptualizations of percepts, derives from their connection to perception. Thus the belief about the tomato by merely sighting the green sliver, or the belief about the square by merely seeing the corner, are warranted through the contribution in present perception by narrower perceptual representations and transitions from them to perceptual belief. *The percepts and transitions from them are the basic sources of empirical entitlement.*³⁶

So it seems that what should be said is this: *Present perceptual representation* makes a contribution to the warrant of a perceptual belief only insofar as the perception involves the perceptual application of representational types whose natures are individuated by types of perceptual paths that connect them reliably to the referents of the perceptual belief. The relevant perceptual belief may or may not be perceptual in the restricted sense of “perceptual”. The warrant-giving paths must go through the perceptions. So perceptual beliefs applying concepts like *car* obtain their perceptual warrant through *present perceptual representation* only through their warrant-yielding association with the restricted set of perceptions—probably by way of the more restricted perceptual beliefs. (Cf. note 36.) Of course, further empirical warrant—which is sometimes essential for being warranted in perceptual belief—can accrue from background empirical beliefs that are brought to the occurrent perceptual situation. But the empirical entitlements and justifications for these, in turn, must ultimately be funded by perceptual states, with their restricted ranges of application.

Entitlement to perceptual beliefs depends on the presence of reason because it applies only within a system of propositional attitudes. One can have propositional attitudes only if some of them sometimes are or have been capable of functioning as reasons.³⁷ Perceptual entitlement can be overturned or supplemented by reasons. But it does not rest on reasons—at least not reasons available to the entitled individual. It is not dependent on control, access, guidance by a rule, much less reflective guidance. The individual need not have the concepts to articulate the warrant. As with deeds of title to property, the conditions that determine whether a candidate entitlement is genuine can depend on matters outside the control and ken of the individual. The reli-

³⁶ I leave open whether the perceptions can contribute directly to the entitlement to the perceptual beliefs in the wider class or whether they contribute to such entitlement only by contributing to the entitlement to perceptual beliefs in the restricted class. In either case, I do not believe that there is an inference from beliefs in the restricted class to beliefs in the wider class. But there are commonly norm-governed relations between the two types of beliefs. These are issues for further empirical and epistemological investigation.

³⁷ Functioning as reasons does not require deliberative self-conscious reflection on reasons, or even understanding the notion of reason. Cf. note 1.

ability and individuation conditions bear on whether a perceptual state is part of a way of forming true perceptual beliefs that meets standards appropriate to the limitations and perspective of the individual.

Entitlement to perceptual beliefs is grounded in two sources. It is grounded in perceptual states with content fashioned through reliable interaction between the environment and sense-perceptual systems. And it is grounded in fulfilling epistemic norms governing competent conceptualization of perceptual representations and formation of perceptual belief. These norms are associated with the representational good, true belief—given the individual's abilities and perspectival limitations. Much of the content and warrant for our rational transactions stems from interaction with a physical world reflected in our animal natures.