

# 2

## Entitlement

### The Basis for Empirical Epistemic Warrant\*

*Tyler Burge*

Most of what we know is supported by beliefs formed fairly directly from sense perception. Many such beliefs constitute knowledge. Over its history, philosophy has not done well in characterizing these beliefs or in saying wherein they are knowledgeable.

One reason is that philosophy has done poorly in characterizing perception. A science of perception was not available until the late nineteenth century, and did not bloom into a rich enterprise until the 1970s. Some of the early mischaracterizations are perhaps to be expected. However, even since the 1970s, philosophy has mostly done poorly in describing perception and in providing an epistemology of empirical belief, because philosophy has not used science well.

Another reason is that philosophy did not focus squarely on ordinary perception-based beliefs until the second half of the twentieth century. The primary topics of epistemological reflection, from Plato through Kant, were more impressive sorts of knowledge. Roughly, the topics were scientific knowledge, including mathematical knowledge, and knowledge gained through reflective understanding. Ordinary sorts of knowledge have not, until the last half-century, been as squarely targeted. The traditional focus on knowledge deluxe is defensible. But it has led to hyper-intellectualized accounts of the first knowledge obtained through sense perception. *That* knowledge is not knowledge deluxe.

A satisfying history of why these distortions occurred has yet to be written. Sometimes it was just assumed that if knowledge closely connected to sense perception is to support rich kinds of empirical knowledge, it must itself be rich, or in some way golden. That assumption is mistaken. Richness accrues from cross-checking, testing, ruling out alternatives, systematization, mathematicization, reflection. Science expands from its humbler beginnings.

Concern with scepticism also fed the distortion. Scepticism is a rich tool for provoking search for philosophical understanding. In the development of ordinary and even scientific knowledge, however, scepticism is a curiosity, not a *rite de passage*. Undermining scepticism is not a requirement on knowledge. Many traditional

\* I am indebted to John Carriero and Barbara Herman for improvements in Section 2.1; to David Chalmers for catching a non-deductive transition in a purportedly deductive argument in an oral presentation of Section 2.2; and to Gary Ebbs and Michael Rescorla for improvements in Section 2.3. I am especially indebted Johannes Burge for valuable comments on several predecessors of Section 2.5.

philosophical accounts of knowledgeable perceptual beliefs build in purported defenses against scepticism, to their detriment.

Whatever the causes, most of the philosophical tradition badly mischaracterizes wherein beliefs most immediately obtained from sense perception are knowledgeable. Much current philosophy still does so. Here I focus on a key normative property that beliefs must have if they are to constitute knowledge. The property is *epistemic warrant*.

The work here comprises six sections. Section 2.1 states some normative principles of epistemic warrant and distinguishes two types—justification and entitlement. *Justification* is warrant through reason. *Entitlement* is warrant without reason. It is the primary type that applies to perceptual beliefs. Section 2.2 discusses the role of reliability in warrant, and the extent to which entitlements to perceptual beliefs depend on factors external to a believer's psychology. Section 2.3 considers some issues about scepticism. Section 2.4 criticizes an argument that believing that we are entitled to perceptual beliefs would commit us to an unacceptable way of validating the reliability of those beliefs. Section 2.5 rebuts an argument that believing that we are entitled to perceptual beliefs is inconsistent with intuitions about confirmation and with Bayesian principles of subjective probability. Section 2.6 concludes.

## 2.1 Epistemic Norms, Justification, and Entitlement

Much work on the epistemology of empirical belief in the last sixty years has focused on its epistemic structure. A modest foundationalism has been a common and salutary upshot. Little of this work has discussed the psychological capacities and the functions and norms for those capacities that make the structure what it is. Form follows function and norms, which follow capacities. In this section, I hope to take some steps toward remedying this situation.

I begin by locating epistemic warrant in a normative structure.<sup>1</sup> *Norms* are standards for success, or for contributing to success, in fulfilling a function, purpose, or goal. *Goods*, in the broad sense relevant here, are successes, or contributions to success, in fulfilling a function, purpose, or goal. Hence goods are fulfillments, or contributions to fulfillments, of norms.

'Epistemic' is etymologically rooted in a Greek word for knowledge, and earlier in an Ionic Greek word for understanding things about the world.<sup>2</sup> Epistemic warrant is an unreducible, normative, epistemic property of instances of propositional psychological states or events—prototypically beliefs, and occurrent thoughts with some

<sup>1</sup> Several of the ideas in the account in this section are present in my 'Perceptual Entitlement', *Philosophy and Phenomenological Research* 67 (2003), 503–548. The account here is more fully developed.

<sup>2</sup> Here, I take 'understanding', in the relevant sense, to be *factive*, like 'knowledge'. Indeed, it is a *type of knowledge* that makes explanatory connections, especially to relatively basic principles. Understanding in this sense should not be identified with minimal competence understanding of representational content. It should not be identified with rich explicative understanding of representational content, though it is closely related. Like ordinary propositional knowledge, understanding here is understanding the world, not understanding representational content. Understanding the world, however, consists in understanding approximately true, relatively deep, propositions or systematic theories.

assertive or at least positive-leaning valence by the thinker. Epistemic warrant can also attach to instances of transitions that lead to such beliefs and thoughts.<sup>3</sup> It can, yet further, attach to instances of reliance on psychological capacities that make use of such beliefs and thoughts—such as reliance on propositional memory or on reflection.

Epistemic warrant is a *normative* property in that it signifies that beliefs, thoughts, inferences, or instances of reliance meet a standard for being good. It is an *epistemic* property in that meeting the standard constitutively contributes toward, or constitutively conduces to, having knowledge.

Epistemic warrant constitutively conduces to having knowledge in at least these respects: (a) A state or event's being epistemically warranted is constitutively necessary for its being knowledge. (b) A state or event's being epistemically warranted supports that state or events' being knowledge. Being epistemically warranted supports having knowledge through that state or event in that, if the epistemically warranted state or event (say, the belief) *is* true, if no Gettier conditions obtain, if there is no internal incoherence, *and if no further epistemic warrant is needed for its being knowledge*, then the belief *is* knowledge.<sup>4</sup> (c) It is impossible, in any case, that a belief's epistemic warrant could block the belief from being knowledge.

In fact, in many cases, epistemic warrant conduces to knowledge in a stronger way. In many cases, epistemic warrant provides at least a *prima facie certification* of knowledge. Constitutively, for some true beliefs, an epistemic warrant certifies knowledge in that it suffices for knowledge, absent counter-warrants, internal incoherence, and Gettier failures. I think that canonical or basic warrants for epistemically basic beliefs certify true beliefs, *prima facie*, as knowledge, in the foregoing sense.

Examples of basic epistemic warrants for epistemically basic beliefs are (i) warrants for perceptual beliefs through forming them directly from perceptual states and (ii) warrants for logical and mathematical beliefs through understanding them.

In sum, basic epistemic warrant for a true perceptual belief certifies it as knowledge, under the stated conditions. In the cases of basic epistemic warrants for beliefs in simple truths of logic and mathematics—that is, warrants through understanding them—being epistemically warranted suffices for knowledge.

The *prima facie* certifying condition does not apply to all cases for warranted true belief. It does not apply to some cases of warranted, statistically based, true beliefs. One can be warranted in a belief on the basis of purely statistical evidence. For example, in a random drawing of balls, nine-tenths of which one knows are red, one is epistemically warranted in believing that the first ball drawn will be red. Even though the belief is true and none of the undermining conditions applies, one may not know that the first ball drawn will be red. In such cases—where the certifying

<sup>3</sup> By 'instances', I intend to signal, for example, instantiations of a psychological state—or event-type in a particular individual at a particular time. One can also think of an *individual's holding* a belief at a particular time as a bearer of epistemic warrant. A given *type* of belief, or a given *type* of inductive inference that has premises and conclusions with the same representational-content types, can be warranted in one individual or time and unwarranted in another. Differences depend on background information available to the individual, or in the level of reliability in the individual's competencies. I explain in more detail below what I mean by a belief *type* or transition *type*. See note 58.

<sup>4</sup> For brief discussion of the lack-of-internal-incoherence condition, see note 82.

condition does not apply—epistemic warrant still constitutively conduces to true belief and to knowledge. It conduces to knowledge in that it meets the conditions (a)–(c) above. In particular, more warrant that would support an explanation of why, in the particular case, a red ball will be formed would make the belief knowledge.

Epistemic warrant always marks a belief as issuing from a competence that conduces to knowledge. Warranted, true beliefs that are based purely on statistics and that do not count as knowledge still conduce to knowledge and to being true.

As I indicated, I think that for every epistemically basic belief, there is a type of possible epistemic warrant that certifies such beliefs, when true, *prima facie*, as knowledge. I also think that some epistemic warrants for some *non-epistemically-basic* beliefs can certify true beliefs of those sorts, *prima facie*, as knowledge. The warranted true beliefs are knowledge, if the stated undermining conditions are not applicable. For example, I think that the relevant certification applies to certain beliefs (such as the belief that is an X-ray machine, applied on the basis of current perception) that are not strictly perceptual beliefs in the sense that I will explicate.

I distinguish between *epistemic warrant* and generic, positive *epistemic support*. Epistemic warrant warrants a belief, at least *prima facie*, as belief-worthy. I think that a belief can have small epistemic support that is not enough to *warrant* belief, even *prima facie*.

In being constitutively conducive to knowledge—specifically, in supporting the relevant state's being knowledge in the sense of (b) above—epistemic warrant constitutively guides toward truth. Being epistemically warranted requires that a belief is formed or held in a way conducive to that type of belief's being true. This conduciveness must reside in a competence in forming or maintaining the belief. The competence must tend, non-accidentally, to connect beliefs of the relevant type, formed in the relevant way, to the subject matter that makes the beliefs true.

These relations to knowledge and to truth are basic to the nature of epistemic warrant. In sum, a psychological state or event is epistemically warranted if it meets certain norms, or standards, for its being epistemically good. They are epistemically good in that beliefs or events of that type, formed or maintained in the relevant way, tend to yield knowledge and hence truth.

Belief and occurrent analogs, such as judgment, are primary psychological states for understanding epistemic warrant—hence for understanding knowledge and understanding. They are primary because their representational function is to be true, and because the central epistemic goods—knowledge, understanding, and epistemic warrant—constitutively serve, in ways to be discussed, representational good or success—specifically, true belief.

In elaborating these points, I start with *representational* goods and functions, as distinct from *epistemic* goods and functions. I focus on functions rather than purposes or goals. I think that beliefs have, by nature, a representational function. They do not literally have purposes, goals, or aims. Talk of the aim of belief is at best metaphorical.

Every belief has the representational function of being true. Any belief undergoes a certain failure *as a belief*—what is naturally called a representational failure—if it is not true. Any belief enjoys a representational success as a belief, if it is true. Success and failure evince functions, goals, aims, or purposes. Beliefs do not themselves have

goals, aims, or purposes. So being true or false is a success or failure in fulfilling a function, a representational function.

Similarly, every perceptual state has the representational function of being accurate. Truth and falsity are properties of propositions and propositional states. Perceptual states are not propositional. So they cannot be true or false. They can be accurate or inaccurate. Any inaccurate perceptual state undergoes a type of failure as a perceptual state. Accuracy is a type of success for perceptual states. Success and failure require presence of functions, goals, aims, or purposes. Perceptual states lack goals, aims, and purposes. So success and failure in perceptual states require functions. The functions are again naturally termed ‘representational functions’.

Fulfillments of functions are goods. So being representationally successful—being true for a belief, being accurate for a perceptual state—is a representational good.<sup>5</sup>

Propositional inference—deductive and inductive—is a primary type of transition for understanding epistemic warrant—hence for understanding knowledge and understanding. Propositional inference undergoes a representational failure if it does not preserve truth.<sup>6</sup> So propositional inference has a representational function of preserving truth. Transitions from perceptual states to perceptual beliefs have the representational function of preserving the veridicality (accuracy) of a perceptual state in the veridicality (truth) of the corresponding perceptual belief.<sup>7</sup> Any such transition that does not preserve veridicality undergoes a type of failure, a representational failure. Propositional inferences and transitions from perception to belief that fulfill their preservational, representational functions are representational goods.

Veridicality of belief and perception is success in fulfilling representational functions. Here is reasoning for taking *representational* functions, hence representational goods, to be distinctive types of functions and goods. As a matter of metaphysical necessity, in every case, beliefs and perceptions fail in some respect if they are not

<sup>5</sup> Of course, I do not think that every true or accurate representational state or content is a representational good. The truth of a proposition about some arcane matter that no one has ever considered is not a good of any kind. Nor is the falsity of its negation a bad thing. Propositions do not in themselves have functions: a proposition does not fail in any way if it is false. I claim that truth and accuracy for certain committal states, like beliefs and perceptions, are representational goods. A *committal* state is one that has a representational function whose fulfillment hinges on whether its representational content is, was, or becomes veridical. Perception, belief, intention, memory, expectation, are committal. Supposing for the sake of argument and playful imagining are non-committal. Whether a modality is committal or not is a priori determinable from the nature of the modality. One cannot renounce the committal character of a state. I take the points in this note, and in the two paragraphs in the text that precede it, to be a priori. Understanding what a belief or perceptual state is requires understanding that it undergoes a certain type of failure, as a belief or perceptual state, if it is not true or accurate—and a certain type of success if it is true or accurate.

<sup>6</sup> Inferences that meet *deductive* standards necessarily fulfill their representational function as a matter of form. From true premises, they necessarily lead to true conclusions by virtue of the inferences’ forms. No inductive inferential transition necessarily fulfills its representational function by virtue of its form. (Inductive inferences in mathematics and other inductive inferences necessarily preserve necessarily true conclusions from necessarily true premises, but not by virtue of the inferences’ forms.) Of course, neither deductive nor inductive inference can insure that premises are true. When representationally successful, both types conditionally preserve truth: if the premises are true, the conclusion is true.

<sup>7</sup> *Veridicality* is the genus. *Truth* and *accuracy* are species. Truth is the propositional form of veridicality. Accuracy is the non-propositional form of veridicality. As noted earlier, beliefs are true or false; perceptual states are accurate or inaccurate.

veridical. They do not necessarily fail with respect to any practical function, if they are not veridical. Any representational state metaphysically could be practically useful but non-veridical, or veridical but practically deleterious. So the function that perception and belief necessarily fail at, if they are not veridical, is not a practical function. Analogously, for relevant transitions. All biological functions are practical functions. They are practical in that they contribute, in some way, to survival long enough to reproduce. So there is a type of function for belief, for perception, and for relevant inferences and transitions that is associated with veridicality and that is not a biological or other practical function. I label this type of function a *representational function*. Meeting the norm for fulfilling this type of function is a *representational good*.<sup>8</sup>

I turn now to connections between representational goods and representational functions, on one hand, and epistemic goods and functions, on the other.

Knowledge, understanding, and epistemic warrant are, of course, epistemic goods. They meet norms for different types of epistemic success. These epistemic norms are constitutively connected to *representational* goods, and to the *representational* function of belief. Knowledge and understanding are epistemically rich ways of *achieving* representational success for a belief—achieving truth. Epistemic warrant *contributes* to a belief's representational success: contributes to fulfilling the master representational norm of belief, being true. A belief's being epistemically warranted is conducive, in an epistemically good way, to the belief's being true.

So all three epistemic goods are good in being epistemically good ways either of a belief's being representationally successful, or in contributing to a belief's representational success. In this very limited sense, true belief is a *telos* of these epistemic goods.

A constitutive condition on knowledge and understanding is that they be true. Any belief undergoes an epistemic failure—a failure to constitute knowledge or understanding—if it is not true. A belief's being true is a constitutive element in central types of epistemic success—in fulfilling some central epistemic norms, in being epistemically good.

True belief is a representational good regardless of how it is produced or maintained. It is not, however, an epistemic good apart from how it is produced or maintained. There is no *epistemic* good in forming a true belief by epistemically bad means.

A belief's being produced or held in a way that is conducive to that type of belief's being true is a constitutive aspect of another epistemic good—epistemic warrant. Epistemic warrant is an epistemic good that is also a constitutive condition on knowledge and understanding.

The epistemic goods, knowledge, understanding, and epistemic warrant, are good in being epistemically good ways of having true belief or, in the case of epistemic warrant, conducing to true belief. True belief is a representational good in itself. True belief is not *in itself* an epistemic good. The epistemic goodness of a true belief hinges partly on the epistemic goodness of how it is formed or maintained. True belief is

<sup>8</sup> The first premise of this argument is set out in 'Perceptual Entitlement', *op. cit.*, 508. The distinction between being true and being useful is also emphasized, 509–510. The argument in the text here, though implicit in the earlier paper, is not explicit there.

epistemically good only through being supported or warranted—or through being knowledge or understanding—only through being formed or held in ways that meet distinctively epistemic norms for being epistemically good.

I stress that, in discussing the epistemic goodness of true belief, I am not discussing what is good for “cognition” as a whole.<sup>9</sup> I have, and have had, nothing to say on goods for cognition in aggregate. I make no claims about how or whether the specific distributed good that I discuss aggregates to make collective goods, or trades off to maximize some overall goodness. *The goods that I specify attach to each belief or inference, one by one.*

In saying that being epistemically warranted is an epistemically good way to obtain truth, I do not imply that epistemic warrant is good only instrumentally, though it is instrumentally good. It meets a standard for having true belief in an epistemically good way—a way that contributes to knowledge and understanding. The epistemic goodness of epistemic warrant lies partly in *how* it conduces to true belief. It must conduce via a competence for forming or sustaining true beliefs of the relevant type, and that competence must be appropriately connected to its successes. Epistemic goodness derives from meeting distinctively epistemic norms. A *telos* of warrant, true belief, is not *epistemically* good independently of beliefs’ being formed or held in epistemically good ways.<sup>10</sup>

Unlike true belief, the epistemic goods, knowledge, understanding, and epistemic warrant, in themselves fulfill epistemic norms. Each of their instances is constitutively an epistemic good. All are epistemically good in being good *ways* of forming, maintaining, or conducing to true belief. True belief is a representational good that these epistemic goods serve.

As noted, inferences are representationally good if they preserve truth. An inference’s preserving truth bears constitutive relations to certain epistemic successes. Any inference, and any transition from perception to belief, incurs an epistemic failure, a failure to yield knowledge, if it does not preserve veridicality. (I do not count

<sup>9</sup> I think that the primary epistemic good of cognition is not to produce beliefs in as many “important” truths as possible while minimizing false beliefs. Here I differ with A. Goldman, *Epistemology and Cognition* (Cambridge, Mass.: Harvard University Press, 1986), 98; and W. Alston, *Beyond “Justification”: Dimensions of Epistemic Evaluation* (Ithaca, New York: Cornell University Press, 2005), 29. In the first place, the way in which one forms or maintains a belief is crucial to its epistemic goodness. In the second, having true beliefs that yield understanding through deep explanations is epistemically good in a way that just having true beliefs is not, even if these latter beliefs are “important”. Even if a belief is warranted, connecting it in explanatory ways to other warranted beliefs is clearly epistemically better than just holding them in a warranted way. (Needless to say, more needs to be said about what counts as being important, on the accounts that I am criticizing.) Connections among one’s beliefs and connections between perceptions and beliefs are crucial to the epistemic good and point of cognition. What is missing in these philosophers’ generalizations are roles for understanding and explanation—and more broadly, a role for specific, epistemically good ways of obtaining true belief, beyond just having true belief.

<sup>10</sup> Truth is a *representational* good for belief independently of how the belief is obtained. I think it obvious that true belief is not an epistemic good independently of how the belief is obtained. For philosophers who may imply that obtaining true belief is *in itself* an epistemic good, see L. BonJour, *The Structure of Empirical Knowledge* (Cambridge, Mass.: Harvard University Press, 1985), 7–8; A. Goldman, ‘The Unity of the Epistemic Virtues’, in *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*, A. Fairweather and L. Zagzebski eds. (Oxford: Oxford University Press, 2001), 31–32.

transitions from perception to belief as *inferences*, because I reserve the term ‘inference’ for propositional transitions that function, representationally, to preserve truth.) Fulfilling the representational function of preserving veridicality can be a constitutive condition on fulfilling certain epistemic norms for epistemic success—norms for having knowledge or understanding.

As with true belief, an inference’s preserving truth is not in itself an epistemic good. There is no epistemic good in a truth-preserving inference that in no way *supports* the truth of its conclusion. To be epistemically good, truth-preserving inferences must meet distinctively epistemic norms.

Analogously to the case of true belief, epistemic goodness in veridicality-preserving inferences and transitions derives from their fulfilling, or contributing to fulfilling, representational functions and norms *in epistemically good ways*. Relevant goods are fulfillments of a norm for fulfilling, or contributing to fulfilling, a function. Epistemic goods, for inferences and perception-to-belief transitions, fulfill or contribute to fulfilling epistemic norms. Doing so serves fulfilling, or contributing to fulfilling, *representational* functions and representational norms of veridicality-preservation.

To summarize, the specific epistemic goods—knowledge, understanding, and epistemic warrant—are good in being epistemically good ways of forming or maintaining true belief. They are good ways of fulfilling or contributing to fulfilling a *representational* function.

Being true fulfills a constitutive representational function of belief. Are there also constitutive epistemic functions for beliefs? In trying to answer this question, I re-use the method used in determining that beliefs constitutively have a representational function. All beliefs constitutively fail as beliefs if they are not true. That fact provides ground to think that all beliefs constitutively have the representational function of being true. A belief constitutively has an *epistemic* function—an epistemic function it has by virtue of being a belief, not by virtue of someone’s using it for some purpose, even an epistemic purpose—if and only if every belief constitutively fails *as a belief* if it does not succeed in fulfilling the function.

Of course, any belief that fails to be understanding or knowledge fails in some respect. It fails to meet an epistemic norm. The issue here is whether it fails *as a belief*—whether it fails some constitutive function of belief.

Not all beliefs fail as beliefs if they do not constitute understanding. Clearly, perceptual beliefs do not constitute understanding. (See note 2.) They do not thereby fail as beliefs. Many other beliefs do not constitute understanding. They do not fall short of norms that apply to all beliefs, as beliefs. Beliefs, *per se*, do not function to constitute understanding.

What of knowledge? I believe that not all beliefs fail as beliefs if they are not knowledge. Beliefs that are not knowledge certainly fail to meet epistemic norms for knowledge. They fall short of realizing a central epistemic good. But failing to be knowledge does not seem to be a defect in a belief *as a belief* on a par with failing to be true. Imagine a person forming a belief about whether a red ball will be extracted blind from a bowl that is known to contain nine-tenths red balls and one tenth green balls. A true belief that a red ball will be extracted is not defective as a belief, if the belief is not knowledge. Similarly, for some beliefs about the future. I believe from weather patterns in Los Angeles that it will not rain in the first week of



October 2023. Suppose that the belief is true. Although it is not knowledge, it is not defective as a belief.

Compare 'I believe that  $p$ , but do not know that  $p$ ' with 'I believe that  $p$ , but  $p$  is not true'. The latter evinces a defect in one's system of beliefs. The former indicates that one's belief that  $p$  does not achieve a certain epistemic good. But it does not necessarily evince a defect in a belief, as a belief. It may involve a wise judgment on the limits of one's perspective. The first-level belief that  $p$  and the second-level judgment that one does not know that  $p$  may even both be epistemically warranted. So being knowledge is not a constitutive epistemic function of belief.

Is being epistemically warranted a function of belief? Perhaps. But perhaps it is possible to have a non-defective belief with some epistemic support, but not enough support to count as warrant. Such a belief might be in a nether region in which it is permissible, perhaps given other considerations, to hold it tentatively. Epistemic warrant is a relatively strong type of epistemic support, solidly supporting belief, not just making it permissible.

However that may be, any belief that lacks epistemic support suffers a failure as a belief, an epistemic failure. Of course an epistemic failure is compatible with a practical success and a representational success. I conclude that it is a constitutive epistemic function of belief to have some epistemic support.

Similar, but more complex, points apply to propositional inferences and transitions from non-propositional to propositional states. In those cases, having an epistemic function hinges on whether the modality of the conclusion of the inference, or the end-point of the transition from a perception, is a belief—and on the point of the inference or transition. An inference that functions to form or maintain a *belief* that provides no support for the belief is defective.

A deductively valid inference that functions to lead to a belief, but is circular, is epistemically defective as an inference. It is epistemically defective even though it fulfills its representational function of conditionally preserving truth. An inductive inference to a belief that preserves truth, but fails to support its conclusion fails as an inference. A transition from a perception to a belief is epistemically defective if the transition does not provide epistemic support for the belief.

Non-defective inferences that function to yield conclusions that are not beliefs but are positive leanings must also yield some epistemic support. Otherwise, they fail in some way, an epistemic way, as inferences. Inferences whose point is not to produce a belief can be non-defective, even if they yield no support. One can make any inference for the fun of it. Transitions from perceptions that do not yield beliefs, but rather withholdings of belief, are non-defective, if, indeed because, the transitions yield no support. An individual could anticipate, because of background knowledge, that a belief based on a coming perception will be false. A transition to a suspension of belief or a dis-belief would not thereby be defective, if the transition yielded no epistemic support for the suspended or rejected belief.

Inductive inferences that provide support for a belief, but that do not preserve truth, fail a constitutive representational function, but fulfill the relevant epistemic function.

Is meeting a norm for knowledge or understanding a constitutive epistemic function of inductive inference? I think not. An argument analogous to the one

about the colored balls that showed that epistemically supported, even warranted, true beliefs need not fail as beliefs, if they are not knowledge, applies. A similar point applies for understanding. Some warranting truth-preserving inductive inferences do not fail *as inferences* by virtue of not yielding knowledge or understanding. Trivially, they fail norms for knowledge and understanding. They do not fail to fulfill a function of all inferences to beliefs.

Our main positive conclusion is that there are epistemic as well as representational functions for beliefs, propositional inferences, and perception-to-belief transitions. Beliefs have a constitutive epistemic function of having epistemic support. Inferences that form or maintain beliefs, or positive leanings, have a constitutive epistemic function of providing epistemic support. Unlike deductive inference, inductive inference to belief can fulfill an epistemic function even if it fails its representational function. Inductive inferences with true premises and false conclusions can yield support, even warrant, for their conclusions. Thus there are constitutive epistemic goods for beliefs, inferences, and transitions from perception to belief.

The toehold for a constitutive epistemic function of belief and relevant transition events grounds a more extensive set of constitutive epistemic functions. For example, it is a constitutive epistemic function of a belief to be maintained at a strength appropriate to epistemic support for it.

Similarly, if an individual forms a belief that is *prima facie* epistemically warranted, and the individual has other accessible information that weighs against that belief, the formation or retention of the belief is epistemically defective if it does not take account of the information, in a way appropriate to its weight.

So far we have discussed belief and inference *per se*. Constitutive functions in these cases must hold for even the most primitive levels of these capacities. Richer representational and epistemic psychological capacities carry richer epistemic functions. A psychology that not only has a capacity for beliefs but also a capacity to form beliefs about beliefs and to evaluate them using epistemic concepts, has wider scope for epistemically evaluable activity and processing. There are, I think, constitutive epistemic functions for different psychological levels of psychological capacity in which belief and inference are embedded.

For example, an individual who can deliberate reflectively about beliefs has epistemically deficient beliefs if deliberation would bear on the epistemic warrant for a first-level belief and the individual does not deliberate. I think that there are epistemic functions and norms that are associated with adequately full use of one's epistemically relevant powers. So I think that there are constitutive epistemic functions, norms, and goods not just for beliefs and inferences *per se*, but for beliefs embedded in specific psychological capacities to carry out explanatory justification.

I do not pursue these matters here. I think that they may present a rich field for better understanding of epistemic functions, norms, and goods.

Wherein is knowledge a good? The goods discussed here are good by virtue of fulfilling, or contributing to fulfilling, a function. I am focusing on representational and epistemic goods, not practical goods. So I lay aside the good of knowledge that derives from someone's wanting it and succeeding in obtaining it or from yielding technology that helps people reach their goals. Being knowledge is not an epistemic function of belief or, I think, any other basic kind of

psychological state or event.<sup>11</sup> So knowledge's being a good is not a result of fulfilling any representational or epistemic function of a basic psychological state or event.

I return to the earlier point that knowledge's being a good depends *partly* on contributing to fulfilling a *representational* function. By fulfilling its epistemic norms, knowledge fulfills the representational function and norm of belief in an epistemically good way. Knowledge is an epistemic good partly in being an epistemically good way of having representational success, a representational good—the representational good of forming or maintaining true belief. Knowledge meets its own norms for contributing to true belief in an epistemically good way. Of course, a belief's being knowledge fulfills not only the belief's representational function. It also fulfills belief's distinctively epistemic function—representing in an epistemically supported way. Obviously, fulfilling these functions does not suffice to fulfill norms for being knowledge.

Similarly, epistemic warrant is a good not only in fulfilling an epistemic function of belief—the function of being epistemically supported. Epistemic warrant is a good also in *contributing* to fulfilling the representational function of belief, by involving conduciveness to the relevant type of belief's being true. Being representationally good or contributing to being representationally good is one constitutive *point* of being epistemically good—of fulfilling epistemic norms—for the central epistemic goods. Being representationally good is representing the world veridically. Being an epistemically good *way* of being or contributing to representational goodness is a distinctive type of goodness.

Knowledge and understanding are such good ways of representing truly that they *over-fulfill* the representational function of belief. They over-fulfill the basic epistemic function of belief—the function of being epistemically supported. One can compare over-fulfillment here with over-fulfilling biological functions. The biological function of a heart is to pump blood well enough to contribute to survival long enough to reproduce. Pumping blood efficiently and healthfully over-fulfills this function. Life has a biological function similar to the function of a heart. Health, vitality, and living long are goods in that they over-fulfill this basic function.

The ways in which knowledge and understanding go beyond true belief make the way in which true belief is formed, or held, especially firmly connected to being true. Knowledge goes beyond true belief partly in being epistemically warranted. An epistemically warranted true belief is formed or held through a psychological competence whose operation is conducive to forming or maintaining true beliefs of the same type. The competence must be constitutively related to the subject matter so as to make the truth of beliefs that are formed or maintained through the competence

<sup>11</sup> I do not count knowledge or understanding as basic psychological kinds, although I leave open whether they are basic in the sense of being undefinable. A basic psychological kind is a psychological kind that falls within psychological laws or law-like generalizations. Psychology does not treat knowledge or understanding as kinds that enter into psychological laws, though obviously they are, in a sense, kinds of psychological state. For more on this matter, see note 55. One might suggest that psychological *systems* have functions of producing knowledge or understanding. I doubt that there is any literal sense in which this is true. Of course, we may aim at knowledge or understanding. Knowledge and understanding are rich goods for both beliefs and psychological systems. And understanding ennobles believers.

deeply non-accidental. Similarly, knowledge goes beyond true belief in avoiding Gettier cases. Without attempting a generalized account of Gettier failures, I think it fair to say that avoiding them necessarily involves one's warrant's being directly connected with the truth of one's true belief. So in both respects—being warranted and avoiding Gettier cases—knowledge involves a way of forming or maintaining true belief that makes connection to truth especially strong. Understanding also over-fulfills the representational function of belief. It over-fulfills it in all the ways that knowledge does. It adds a systematizing, explanatory epistemic value.

Of course, it does not follow from any of this that the epistemic reduces to the representational. Meeting an epistemic norm does not reduce to fulfilling, or contributing to fulfilling, a representational function or representational norm. The point is, rather, that epistemic norms, in central cases, are standards for distinctively epistemic ways of being representationally good. The central epistemic goods are deeper and richer than representational goods.

The reasoning about functions, norms, and goods in this section has been apriori. I think it interesting and valuable to have discovered, apriori, some basic representational and epistemic functions, norms, and goods for belief and for those inferences or other transitions that function to yield belief. The goods do not depend on anyone's caring about them or pursuing them. They attach apriori and necessarily to relevant psychological capacities.

I focus now on epistemic warrant, beginning with an overview of its fundamental aspects.

Epistemic warrant fulfills a norm for an individual's, or a psychology's, operating competently and in a way that conduces to having knowledge, *given* the information available to the individual. Epistemic warrant attaches paradigmatically to beliefs and occurrent committal thoughts—relative to ways of forming or maintaining them—to inferences to committal conclusions, and to reliance on capacities that make use of beliefs or inferences. (See note 5.)

The epistemic goods, knowledge and understanding, are factive. They constitutively involve true belief. Epistemic warrant is non-factive. A belief can be warranted without being true. An inference can be warranted even though it leads from true premises to a false conclusion.

The distinctive aspect of epistemic warrant that allows a false belief to be warranted is that it allows for limitations on information available to individuals. For example, an individual's false perceptual belief can be warranted if the individual's belief falls into brute error. *Brute error* is error that derives from a well-working psychological system, such as a perceptual or belief-forming system, that happens to be in a situation that yields the error.<sup>12</sup> The error is not the fault of the system or individual. Every belief-forming system that relies on perception is subject to brute error. It can operate as well as possible, given the information that it has, and still yield a false belief. Similarly, a warranted inductive inference from true warranted premises may support a warranted but false conclusion. To be warranted, the inference type must tend, in normal conditions, to yield true belief given true

<sup>12</sup> 'Perceptual Entitlement', *op. cit.*, 507.

premises. But in the given case, there may be a factor that the individual lacks information about and that results in the inference's being unsound.

A second aspect of epistemic warrant is that it meets norms for operating so as to contribute to gaining knowledge. This second aspect is codified in conditions (a)–(c) above.

Epistemic warrant's contributions to knowledge are what make epistemic warrant an *epistemic* property. It contributes to knowledge by conducing to knowledge—(a)–(c). More strongly, as noted, epistemic warrant contributes by *certifying* certain types of true beliefs, *prima facie*, as knowledge. In the absence of countervailing warrants, internal incoherence, and Gettier failures, epistemically basic true beliefs, warranted in epistemically basic ways, *are* knowledge.

A third aspect of norms for epistemic warrant is that they mark operating well in forming or maintaining *true* belief or *preserving truth* in inference. To operate relevantly *well*, the operation must tend toward forming true belief or preserving truth in inference, given the information available. If the relevant propensity to true belief is to contribute to knowledge, as epistemic warrant does, it must derive from the psychological competence. The formation of belief, whether inferential or not, must be part of a non-accidental pattern—capable of grounding a general law-like explanation—that tends to connect the belief-forming competence itself with the subject matter that makes beliefs of that type true.<sup>13</sup> Epistemic warrant is not factive. But these connections between epistemic warrant and conduciveness to truth are just as constitutive of epistemic warrant as the factive connections are to knowledge and understanding.<sup>14</sup>

<sup>13</sup> I emphasize this point in section VI of 'Perceptual Entitlement', *op. cit.*. See further discussion in Section 2.2.

<sup>14</sup> For an attempt to show that there is something wrong with taking epistemic warrant to have a teleological character, conduciveness to truth, see S. Berker, 'Epistemic Teleology and the Separateness of Propositions', *Philosophical Review* 122 (2013), 337–393; and 'The Rejection of Epistemic Consequentialism', *Philosophical Issues* 23 (2013), 363–87. The arguments hinge on taking, indeed defining, teleological approaches in epistemology so that they bear structural relations to simplistic, unattractive consequentialist ethical theories. The arguments are strong and insightful against the positions that they target. Some philosophers have written things that suggest the specific type of teleological position that the arguments target. But it is not *clear* to me that the philosophers whom Berker criticizes believe, reflectively, that being conducive to truth is a value for warrant and "cognition" in the way that his arguments assume. Specifically, the normal reading of 'an epistemic warrant is conducive to true belief' is not that an epistemic warrant is conducive to true belief in general, regardless of what truth is at issue. The normal reading is that an epistemic warrant is conducive to the truth of beliefs of the type that the warrant warrants believing. Berker's arguments hinge on opponents' taking warrant's conduciveness to truth in the former way rather than the latter. In any case, I am committed to none of the tenets of the teleological theories that he imputes to the philosophers whom he criticizes. I do not take true belief to be a final value in his sense. I propose no theory of overall value. I do not accept a deontic theory of the sort that his arguments require. I do not see an epistemic warrant as "promoting" truth in general, as opposed to marking a good way of forming or sustaining the relevant sort of belief. Still, contrary to his definition of a 'teleological' view, my view has teleological commitments. Epistemic warrant marks a way of producing or sustaining the belief as meeting standards for being an epistemically good way of meeting the belief's representational function. Berker presents his arguments as if they defeat all teleological conceptions of epistemic warrant (or "epistemic justification"—see note 17). He recommends giving up any requirement that a warrant conduce to truth. He gestures toward a guidance conception criticized below in Section 2.2. I believe that this position is conceptually wayward. Understanding what epistemic warrant is requires that one take it as marking conduciveness to both knowledge and true belief. As argued in Sections 2.1–2.2, it marks a certain sort of reliable true belief.

Although there are issues about when a competence can include a prosthetic device and remain the individual's competence, I shall not investigate these issues here. I take it that perception through glasses counts as part of one's competence, but a benevolent scientist's hooking one up unawares to a special truth-finding machine cannot be included in one's competence. The issues that I discuss do not hinge on how to adjudicate hard cases for determining what an individual's competence consists in.

An individual's holding an epistemically warranted belief can be conducive to the belief's being true because of the nature of the belief itself, as in the case of an individual's acceptance of obvious logical or mathematical truths through understanding them. Or more commonly, it can be because of how the belief is related to other psychological states and epistemically relevant conditions. An epistemically warranted belief is situated in repeatable psychological patterns. Relevant patterns include competent production of beliefs of the relevant type. Those patterns are part of larger patterns that involve relations of the competence to the belief's subject matter. The larger patterns conduce to beliefs of the relevant type being true, through the relevant competence's relations to the belief's subject matter. I sloganize this point by saying that to be epistemically warranted, a psychological state or event must conduce to truth by being *on a good route to truth*. The route must involve a competence that is itself non-accidentally connected to forming true beliefs.

Deductive inferences preserve truth by their form and nature. Warranted inductive inferences are conducive to preserving truth through their constituting psychological patterns that mirror and derive from environmental patterns. The environmental patterns tend to connect attributes represented by premises with attributes represented by conclusions. Warranted inductive inferences form parts of epistemically good routes to truth.

Epistemically warranted belief must be on a route to truth that is relevantly *reliable*. Constitutively, epistemic warrant marks competencies for producing or sustaining belief as relevantly reliable in producing or sustaining true belief. This requirement derives from epistemic warrant's role in marking a belief or inference as being conducive to knowledge, specifically (c). The formation of a belief or the carrying out of an inference that is not relevantly reliable *could not* be conducive to knowledge. No competence that formed or sustained a belief, or was involved in a belief-supporting inference, could be good enough to be conducive to knowledge if it were in every relevant way unreliable in yielding true belief or in preserving truth in inference. Indeed, if being epistemically warranted (*per impossibile*) allowed a competence that formed or sustained a belief to be relevantly unreliable, then being epistemically warranted could mark a competence that was a decisive hindrance to having knowledge—contrary to condition (c) on being conducive to knowledge. No belief that relied on the competence, even partially, to meet conditions on knowledge could possibly be knowledge.

As indicated, I think that in certain cases—including epistemically basic types of belief—canonical or basic epistemic warrant for those beliefs *certifies* true beliefs *prima facie* as knowledge. In the absence of stronger counter-warrants, of internal incoherence, and of Gettier failures, epistemic warrant marks many true beliefs as knowledge. In particular, under such conditions, the basic type of epistemic warrant for perceptual belief certifies true perceptual beliefs *prima facie* as knowledge. Again, such certification requires that the competence that supports the true belief be relevantly reliable.

The most general point is this: Since, in *all* cases, epistemic warrant marks beliefs and inferences as being conducive to knowledge—as a necessary condition of being an *epistemic* warrant—epistemic warrant marks beliefs and inferences as formed or maintained in a reliably veridical way. I return to the *relevant* reliability, for empirically based beliefs, in Section 2.2.

Being epistemically warranted also requires that the warranted state or event derive from a competence that *functions well* in the given case, in forming or maintaining the belief.<sup>15</sup> The standards that epistemic warrant meets concern an individual's psychological competencies' contribution to getting things right in a way conducive to knowledge. To be part of a good route to truth that derives from a competence's being connected to a belief's subject matter, the *competence*, as distinguished from a malfunction or an intrusion from the outside, must be the source of the belief. If a malfunction or intrusion figures essentially in producing or maintaining a belief, the competence is not the main psychological source. The exercise of the competence does not meet the standard of operating well in serving knowledge. The exercise of a well-functioning competence is the individual's psychology's contribution to being epistemically warranted.

A further reason for taking a well-functioning competence to be a condition on epistemic warrant is the following: The norms for epistemic warrant are standards for forming or maintaining beliefs, or engaging in inferences, in ways that are conducive to gaining knowledge. A true belief that resulted from a malfunction of or intrusion onto the competence could never be knowledge, because of condition (c), even if the malfunction or intrusion reliably produced true belief.

In 'Perceptual Entitlement', I distinguish two types of epistemic warrant—*justification* and *entitlement*.<sup>16</sup>

A *justification* is a warrant whose force derives at least partly from a reason. A psychological *state* or *event* is *justified* if and only if it is epistemically warranted, and the force of its epistemic warrant derives, at least partly, from a reason (or reasons) in the psychology.<sup>17</sup> An *individual* is justified regarding a given state or

<sup>15</sup> I get this point wrong in 'Perceptual Entitlement', *op. cit.*, 537n24.

<sup>16</sup> 'Perceptual Entitlement', *op. cit.*, section I.

<sup>17</sup> Many philosophers use 'justification' roughly as I use 'warrant'. I use 'justification' as I do because the term is traditionally and etymologically tied to 'justify', which is to *show* something to be just or right, or to provide an acceptable explanation for something. Entitlements concern a certain status that does not in itself involve any such showing or explaining. Hence they are not naturally called justifications. Entitlements commonly attach to perceptual beliefs. Such beliefs need not be supported by argument. I think that use of 'justification' (and German cognates) goes fairly far back in the history of philosophy—it is certainly present in Kant. But use of the term to cover beliefs that are not "justified" by further reasons or by self-evidence, but rather by some default status, is more recent and less historically entrenched. Early influences in this unfortunate usage are articles by W. Alston from the 1970s and early 1980s, collected in his *Epistemic Justification* (Ithaca, Cornell University Press, 1989).

Traditionally and etymologically, a warrant is a kind of pledge or guarantee. A warrant may or may not involve a showing or an explanation. 'Warrant' is better fitted to the broader category than 'justification' is. Incidentally, I hope that it is obvious that my use of 'warrant' differs from the use in A. Plantinga, *Warrant: The Current Debate*, *op. cit.* (Oxford: Oxford University Press, 1993). I think that Plantinga's use, which entails that an individual is not epistemically warranted if Gettier conditions are not satisfied, is idiosyncratic.

Some philosophers use 'reason' for any warranting factor. A pain, a perception, or even an environmental condition can, on that usage, count as a reason. Perceptual beliefs are counted 'reasonable' and 'reason-based'. This usage perniciously blurs distinctions among psychological and normative kinds.

event in his or her psychology, if the individual, or the psychology, connects warranting reasons to that state or event. A reason for a conclusion is a propositional state, actual or possible, that provides support for the conclusion through actual or possible reasoning from the state's content to the conclusion.

A reason can provide either *prima facie* or *all-things-considered* justification. It provides *prima facie justification*, if it provides some justification, and would provide sufficient justification for belief, inference, or what not, if no counter-warrants were involved. It provides *all-things-considered justification*, if in the actual case, including all relevant factors, it provides sufficient justification for belief, inference, or what not.

I follow standard practice in distinguishing an individual's *having* a justification for a state or event from that state or event's *being justified*. A justification can be present in or easily available to an individual's psychology, even if the justification is not operative in being appropriately connected in the psychology with what it is reason for. Then the individual *has* the justification. For the state or event to be justified, the individual must not only have the justification. That justification must be *operative*. It must be connected to the state or event in such a way as to support it in actual reasoning or in a psychological structure.

I assume that reasons are propositional. They form part of a propositional sequence that in effect constitutes explanatory reasoning as to why the state or event that they are reasons for is epistemically good. In effect, reasons jointly answer possible normative why-questions with normative because-answers.<sup>18</sup> Perceptual states lack propositional structure. They can be cited in reasons. They can figure constitutively in an entitlement. They cannot literally be steps in reasoning that explains anything. So they cannot themselves be reasons.

Reasons need not involve a capacity for reflection. Reasoners need not have any epistemic or other meta-representational concepts. Hence reasoners need not have a capacity for self-consciousness. Propositional inferences by very young children who lack reflective capacities and consciousness of self can be reasoning and can yield justifications. Reasons need not even be accessible to consciousness. A structure or operation's meeting rational norms does not require that it be accessible to consciousness. *Having* a reason requires only that the reason occur in the individual's psychology—or be easily inferable, given the individual's psychological powers, from what *is* in that psychology.<sup>19</sup> An individual's belief's *being justified* requires

[Note 17 continued:]

I hope to discuss such issues elsewhere. I stress here that I take reasons to be propositional representational contents, or possible attitudes with such contents.

<sup>18</sup> I say 'in effect' because no one need pose a question about why a conclusion is belief-worthy. A reasoner need not even *be able* to pose such a question. There is empirical reason to think that children reason before they have a concept of belief and can pose relevant questions. Their reasons nevertheless support the conclusion in a way that could ground an explanation of its belief-worthiness.

<sup>19</sup> Some philosophers engage in a usage on which an individual "has" any reason that is applicable to the individual's situation, whether or not the reason is in the individual's psychology or otherwise available to the individual. See T. Scanlon, *What We Owe Each Other* (Cambridge, Mass.: Harvard University Press, 1998); *Being Realistic About Reasons* (Oxford: Oxford University Press, 2014). On this usage, an individual can "have" a reason even if he or she lacks concepts needed to think the content of the reason. I regard this



only that the believer have reasons, and that they operatively support, in the believer's psychology, what they are reasons for.<sup>20</sup> As noted earlier, some epistemic support for

usage as linguistically odd. I also think it detached from epistemology. It detaches reasons from potential reasoning by the individual who "has" the reasons. I think of reasons that are applicable to the individual's situation, but that the individual does not have in his or her psychology, as potential reasons, or God's-eye reasons. I hope to discuss these matters elsewhere.

<sup>20</sup> Here some remarks on the contrast between this characterization of justification and the one in 'Perceptual Entitlement', *op. cit.* are in order. The main contrast lies in a different position on accessibility. In 'Perceptual Entitlement', 505, I characterized *justification* as 'warrant by reason that is conceptually accessible on reflection to the warranted individual'. I stressed (505, 519–521) that reflection is not necessary for having reasons. So I meant 'is conceptually accessible, if reflection could be and were to be exercised'. In explicating *entitlement*, 504, I wrote that it is warrant that is not accessible *even* on reflection. Reflection is an especially strong tool for bringing something to consciousness. The relevant necessary condition for justification was simply a requirement that the individual be able—have a competence—consciously to think a propositional formulation of the warrant (504). For justification, the warrant consists in the warranting reasons together, in most cases, with an inferential transition from the reasons to a conclusion. So the access requirement on justification was merely that the individual be able consciously to think the reason (not necessarily think it as a reason). I emphasize that, even in 'Perceptual Entitlement', my access requirement on justifications was very spare—not at all demanding. To be justified, the individual also had to have a disposition to connect the reason, for example by inference, to the conclusion—an inference that *could* be made conscious. I contrasted justification with entitlement. I claimed that the former but not the latter requires conceptual access to the "warrant". The former but not the latter required an ability to bring the warrant to conceptual consciousness.

I later dropped the conscious accessibility requirement on justification. See 'Epistemic Warrant: Humans and Computers' (2011), in my *Cognition Through Understanding: Philosophical Essays, Volume III* (Oxford: Oxford University Press, 2013). I thought in 2003 that reasons must be accessible to being consciously thought. In 2003, I glossed the lack of access-requirement in the case of entitlement: 'The individual need not have the concepts necessary to *think the propositional content that formulates the warrant*' (504, italics added here). The warrant in the case of perceptual entitlement is the entitlement. For perceptual belief, the entitlement is the process of belief formation from registration of proximal stimulation through a perceptual state to a perceptual belief. To be an entitlement, such a process must reliably produce accurate perceptions and true beliefs generated from the perceptions; the reliability must be grounded in relations to the environment that help explain the connection between reliability and the warranted beliefs; and the process must be psychologically well-functioning in the particular instance. In fact, I believed, and continue to believe, that to be entitled to a perceptual belief, an individual need not be able, consciously or unconsciously, to formulate or specify a *single aspect* of the entitlement.

Some commentators did not think seriously about this passage ('Perceptual Entitlement', 504). They produced badly mistaken glosses on my access requirement and on other aspects of my justification/entitlement distinction. They imported notions into their explications that I nowhere used. And they overread passages, or read them out of context. See A. Casullo, 'What Is Entitlement?', *Acta Analytica* 22 (2007), 267–279 and N. Silins 'Explaining Perceptual Entitlement', *Erkenntnis* 76 (2012), 243–261. A largely sound reply to Casullo occurs in B. Majors, 'What Entitlement Is', *Acta Analytica*, 30 (2015), 363–387. I hope to criticize these readings in more detail on another occasion.

In 'Perceptual Entitlement', I used a conscious-access requirement on *justification*, because I thought that all reasons could in principle come to consciousness. I thought that warranting factors in an entitlement need not ever be capable of being consciously thought, or even conceptualized. My claim that an individual need not be able to conceptualize the individual's own entitlements entailed that entitlement was warrant without reason. And I thought it obvious that much, or all, of the entitlement for perceptual beliefs was never consciously accessible. So consciousness was not a key issue for understanding entitlement. Subsequently, I realized that there is nothing apriori impossible about reasons in an individual's psychology that are not *consciously* thinkable by the individual—in effect, modular reasons. Since I regard the distinction between warrants by reason and warrants without reason as fundamental, I dropped the access requirement on justifications. The distinction between justification and entitlement is now independent of issues about consciousness. See 'Epistemic Warrant: Humans and Computers', *op. cit.* Of course, the accessible-reason/inaccessible-reason distinction remains philosophically important.

a belief is too weak to constitute warrant or justification. The support may support only a slight favorable leaning. So not all reasons yield even *prima facie* justification. Operative support can be via actual inference or via structurally sustaining a reason's conclusion.

Since knowledge is propositional and perceptual states are not, perceptual states cannot be knowledge. Nor can memory that simply preserves perceptual representation. Since epistemic warrant is a property of states or events that can constitute knowledge, perceptual states and perceptual memories cannot have epistemic warrant.<sup>21</sup> They can have an *analog* of epistemic warrant. By being reliable and the product of a well-functioning perceptual system, they meet an epistemically relevant norm for helping to form warranted perceptual beliefs.

A psychological state, event, or transition has an epistemic *entitlement* if it is epistemically warranted and the force of its epistemic warrant does not derive, even in part, from a reason.<sup>22</sup> A perceptual belief is warranted inasmuch as it fulfills the norm for being warranted. What fulfills the norm is a sequence of event types that begin with a sensory registration, lead to a perceptual state, and thence to the perceptual belief. Fulfilling the norm is meeting conditions that I will discuss. An *individual* is entitled to a belief or inference if the individual has an operative epistemic entitlement to it.

Entitlement is warrant without reason.<sup>23</sup> An entitlement can always in principle be paralleled by a counterpart justification. A sufficiently sophisticated and knowledgeable individual could specify the entitlement in propositional form. The force of an entitlement does not, however, derive from any such specification. An individual whose state is warranted by entitlement need not be *capable* of thinking 'the propositional content that formulates the warrant', by thinking counterpart reasons.

An entitlement derives from the epistemic goodness of the competent production of the warranted state or transition, where the competence is not that of producing a reason. The competence may be that of forming a reliably accurate perceptual state and transitioning reliably from it to a reliably true perceptual belief. Epistemic goodness resides in the warranted state's being produced in a well-functioning, relevantly reliable competence. Most of our perceptual beliefs are warranted by

<sup>21</sup> For further discussion of the point that epistemic warrant is a feature only of propositional states, which do not include perceptual states, see 'Perceptual Entitlement', *op. cit.*, 516–518, 530–533, esp. 533, 535, 540n, 544n.

<sup>22</sup> In 'Perceptual Entitlement', *op. cit.*, 504, my primary explication of entitlement states that a warrant is an entitlement if an individual's having the warrant does not require that the warrant be accessible to the individual or that the individual even be able to think any propositional content that formulates the warrant. Since I explicitly indicated that I conceived of reasons as constitutively propositional (for example, 528), it follows from the older explication that entitlements are warrants without reasons. Entitlements are warrants that do not derive from reasons that the individual has. In my view, the explication of entitlement in 'Perceptual Entitlement' (in contrast to the explication of justification—see note 20 above) is extensionally equivalent to the newer explication. The present explication of justification and entitlement is stated in 'Epistemic Warrant: Humans and Computers', *op. cit.*, 490.

<sup>23</sup> Although my explications for justification and entitlement are for *epistemic* warrant, there are parallel notions for practical warrant: practical justification and practical entitlement. For work on practical warrants analogous to epistemic entitlements, see Y. Luthra, *Embodied Rational Agency*, UCLA Dissertation, (2013), chapter 6—retrieved from <https://escholarship.org/uc/item/4q92g4w6>.

entitlement. To be warranted, they need not be backed by any reason. Similarly, the relevant competence may be that of forming a truth-preserving deductive inference or a truth-conducive inductive inference. Most of our inferential transitions, deductive and inductive, have entitlements, not justifications, at least until we attain a considerable level of sophistication. Giving reasons for the goodness of an inferential transition—perhaps showing that an inference rule is a good one—is a sophisticated enterprise.

My discussion centers on the epistemic status of perceptual beliefs. Here is a good place to explain what I mean by ‘perceptual belief’. *Perceptual belief* is belief formed directly from perceptual states. Being formed directly from perceptual states entails meeting four conditions.

First, the formation must be not be through (propositional) inference.

Second, all elements that function to represent in a *de re* way in a perceptual belief content depend for fulfilling their representational function on deriving from counterpart elements that function to be *de re* in the representational content of the perceptual state that underlies the perceptual belief.<sup>24</sup> I call representational, context-bound, referential elements that function to be *de re* in belief contents and other propositional contents, ‘*referential applications in propositional thought*’. I call representational referential elements that function to be *de re* in perception, ‘*referential applications in perception*’. In what follows, I often shorten ‘referential applications’ to ‘applications’. Applications that function to be *de re* (say, applications of singular-demonstrative-like representational elements) are applications that function to pick out a *representatum* in a non-inferential, not-completely-conceptual-or-perceptual-attributitional, context-dependent way. Applications of either type are individuated via events at a given time. They are semantically like actual occurrent uses of demonstratives. Applications in perception are distinguished from applications in propositional thought in that they are referential applications of perceptual attributives, rather than conceptual attributives. *Perceptual attributives* are attributives formed in a perceptual system. All applications in a perceptual belief that function to be *de re* depend for successful *de re* representation on corresponding successful applications in the perceptual state.

Third, every concept in a perceptual belief *conceptualizes* a perceptual attributive in the underlying perceptual state.<sup>25</sup> Usually, *not all* such perceptual attributives are conceptualized. But every conceptual attributive in a perceptual belief must have a counterpart perceptual attributive in the perceptual state from which the belief (and

<sup>24</sup> A *de re* psychological state is (roughly) one that functions to pick out a *representatum* in a direct, non-inferential way, and that does not depend fully on concepts or perceptual attributives for doing so. For perceptual states and perceptual beliefs, *de re* representation is representation that involves occurrent demonstrative—or indexical-like elements. For less rough discussion, see my ‘Postscript to “Belief De Re”’ *Foundations of Mind* (Oxford: Oxford University Press, 2007).

<sup>25</sup> For discussion of conceptualization of perceptions, see ‘Perceptual Entitlement’, *op. cit.*, section VIII. For discussion of conceptual and perceptual indication and attribution, see my *Origins of Objectivity* (Oxford: Oxford University Press, 2010), 30–38; and ‘Predication and Truth’, *The Journal of Philosophy*, 104 (2007), pp. 580–608. Determining what attributives are perceptual is a delicate empirical matter. See my ‘Reply to Block: Adaptation and the Upper Border of Perception’, *Philosophy and Phenomenological Research* 89 (2014), 573–583. Attributives for color, shape, motion, and body are clearly produced in visual perceptual systems. It is empirically plausible, but less clear, that hand is a perceptual attributive for

the conceptual attributive) is formed. The concept must indicate and attribute the same attribute that the perceptual attributive does, and must depend essentially on the perceptual attributive for doing so.

I emphasize that conceptualizing perceptual representations is not representing them or being about them. It is producing concepts that indicate and attribute the same environmental entities that corresponding perceptual attributives in the perceptual state do, while making essential use of those perceptual modes of presentation to do so.

Here is a concrete illustration of this third condition. A visual perception as of a body of a certain shape and size will figure in the formation of a visual-perceptual belief as of the body of that shape and size. The visual perceptual state has attributives for *body* and for a specific shape and specific size. The perceptual belief in which these perceptual attributives are conceptualized has counterpart concepts for *body* and for the same specific shape and specific size. The perceptual system will commonly also have more generic attributives for shape and size; and these attributives may also have counterpart perceptual concepts in perceptual beliefs. The representational content of the perceptual belief is partly made up of conceptual counterparts of visual perceptual attributives that indicate and attribute some of the same environmental attributes that the perceptual attributives do. The conceptualizations of perceptual-attributive counterparts represent by depending on representation by the underlying perceptual attributives. The conceptualizations differ from the perceptual attributives primarily in that they function to contribute to a propositional structure. No concept in a perceptual belief lacks a counterpart attributive in perception that indicates the same attribute that the concept does.

So if a perceptual system produces no attributive that represents x-ray machines as such, there are no perceptual beliefs—in my sense—that represent bodies as x-ray machines. It is a substantive empirical question which beliefs are perceptual beliefs, in my sense. For it is an empirical question what attributives are formed in perceptual systems.

Fourth, every attributive that is referentially applied in the perception and that is conceptualized in the belief is referentially applied in the belief. The perceptual belief has *de re* applications that match *de re* applications in the conceptualized perception.

Imagine having a perception as of a red body. Imagine that one perceives the body and the instance of redness. There are perceptual applications both to the body and to the instance of redness. In English, we express what we colloquially call a ‘perceptual belief’ formed from the perception as a belief that<sub>1</sub> body is red. It is understood that the subscript on ‘that’ marks a referential application, by or in the believer, of the attributive body that functions to refer to a particular body. However, the English here does not express what I count as a perceptual belief. It does not express the *de re* application purportedly to an instance of redness. The belief as expressed here could be true if one saw the body, but did not see its redness. Perhaps the body is too distant to discern its color. One might believe correctly that it is red through interlocution, or

primates. I count it one in this chapter. I think it unlikely that mule—as opposed to animate body of such and such shape, where the shape is in fact mule-like—, playing card, x-ray machine, or particle accelerator is a perceptual attributive.

on a whim. Then the belief would be a hybrid of perceptually dependent elements and elements from interlocution or whim. It would be a hybrid even if the conceptual attributive red conceptualizes a perceptual attributive red that had occurred in previous perceptions.

To be a perceptual belief, in my sense, the conceptual attributive red must be applied in a *de re* way—given that the counterpart perceptual attributive red is applied *de re*. A way of expressing a relevant perceptual belief in English would be that<sub>1</sub> body is red<sub>that2</sub>. Both the concept body and the concept red must be accompanied by applications that function to apply *de re* to particulars. So every concept in the perceptual belief not only has a counterpart perceptual attributive. Every concept in the perceptual belief that conceptualizes a *de re* applied perceptual attributive must also be *de re* applied in the belief, in a way that depends on the counterpart *de re* application in perception.<sup>26</sup>

The foregoing explication of ‘perceptual belief’ is strict. One might prefer to use the term ‘basic perceptual belief’ or ‘purely perception-dependent belief’. My motive for centering on this notion is to isolate those beliefs that depend on perception as directly as possible. I think that all empirically supported beliefs either are or are supported by perceptual beliefs in this sense.

Perceptual beliefs are not formed by propositional inference. Nor are they formed from non-inferentially associating some perceptually represented attribute with a *further* attribute that is not indicated or attributed in the perceptual system. Such extensions go beyond converting perceptual representation into perceptual concepts and perceptual beliefs. They involve concepts that are not strictly perceptual concepts. A further system beyond that of conceptualizing perceptual attributives, such as a system of associative learning in long term memory, is involved. If visual perceptual systems do not produce attributives that represent x-ray machines as such, there are no perceptual beliefs, in my sense, that represent x-ray machines as such.

By the third condition, a perceptual belief indicates and attributes no attributes not already indicated and attributed in the underlying perception. By the second condition, a perceptual belief makes no applications in thought that do not depend anaphorically on applications in the underlying perception. By the fourth condition, each *de re* application of a perceptual attributive that is conceptualized in the perceptual belief has a counterpart *de re* application in the perceptual belief. So perceptual belief can omit perceptual attributions and applications, but it does not generalize on, abstract from, or add to them. Such generalization, abstraction, and addition are not functions of perceptual belief. They have their own epistemic norms.

We do *say* that the expert sees the object as an x-ray machine. We say this even if there is no perceptual attributive that applies specifically to x-ray machines as such.

<sup>26</sup> Some perceptual attributives are not referentially applied in occurrent perceptual states. These are perceptual attributives for relations, like spatial relations, whose instances lack relevant causal powers and are not, strictly speaking, perceived, although the relation types are perceptually attributed. One does not strictly see the spatial relation between two entities that one perceives as next to one another. One sees concrete particulars *as* next to one another. In such cases, the conceptualizations, in a perceptual belief, of the relational perceptual attributives need not accompany a referential application to an instance of the relational attributive in the perceptual belief. Such relational perceptual attributives always function, however, to relate particulars, the *relata*, that are picked out by applications, hence perceived.

Such talk may seem to be in conflict with the third condition. But we should not take the talk to imply a judgment about whether the expert's perceptual system itself produces representations as of x-ray machines. What attributives are produced in a perceptual system are for empirical science to determine. An expert's association might be so automatic and non-inferential that it is colloquially natural to call the result a 'perceptual belief'. I am interested in norms governing actual mental activity and actual psychological natural kinds—not what ordinary language allows us to say. A fair amount is known about what attributives perceptual systems produce. Epistemic norms are fitted to our specific psychological capacities. I want a term that tracks perceptual beliefs that are as close to perception as possible. I reserve the term 'perceptual belief' for this narrow range of beliefs. I focus on entitlements to beliefs most closely derived from perception. Such focus not only limits complexity. It also tracks epistemic entitlements that are the basis for all empirical warrant.

Most of what I say about entitlement for perceptual beliefs applies, however, with minor adjustments, to a wider class of beliefs that are colloquially called 'perceptual'. The attributives in the beliefs in this wider class need not have counterpart perceptual attributives that indicate the same attributes—as long as they are non-inferentially, *de re* applied by way of perception.

For example, suppose again that there are no perceptual attributives—attributives produced in a perceptual system—that indicate x-ray machines as such. That is, suppose that no perceptual system produces attributives specifically as of x-ray machines. Suppose that an individual's psychology contains non-inferential connections between applications of perceptual shape-, size-, and body-attributives and applications of a conceptual x-ray-machine attributive. Suppose that if the individual has a perception as of a body of the relevant shape and size, the individual non-inferentially forms a defeasible belief that<sub>1</sub> body is an x-ray machine<sub>that2</sub>. (Note that the attributive function is an x-ray machine, marked by the copula, outranks in scope the second referential application.) Suppose that there is a strong associative relation between a conceptual attributive x-ray machine and a perceptual concept as of a body with a certain shape and size. The association is stored in long-term conceptual memory. So the referential application of x-ray machine finds its referent via perception as of a relevant body, shape, and size. The subscript on 'x-ray machine' indicates that although the conceptual attributive x-ray machine is the main predicate of the propositional content, it functions *de re* to betoken (roughly, refer to) a particular x-ray machine. It may or may not succeed in referring. Reference here is not infallible, since perceptual reference is not infallible. Suppose that the x-ray machine is *perceived*, though not strictly *perceived as* an x-ray machine. The conceptual attributive is applied through its non-inferential associative connection to a perception of and as of the body, where the body constitutes the x-ray machine. A condition on successful application of that<sub>2</sub> to an x-ray machine is that the application that<sub>1</sub> refers to the body. Successful application of x-ray machine in thought depends on successful application in the perceptual system of a *perceptual* attributive body.

One can be *entitled* to the belief that<sub>1</sub> body is an x-ray machine<sub>that2</sub>. The association relation is not one of propositional reason support. One's warrant is not a reason. Reasons at least sketch explanations of the belief-worthiness of a conclusion.

Association is not explanation. Even where there is a transition from the perceptual belief as of a body and its size and shape to another propositional attitude, the x-ray-machine belief, the transition need not be one of propositional inference or reason support.<sup>27</sup> The formation of the belief might be warranted, if it derives from a reliable, well-functioning, appropriately grounded, associational competence. So the not-strictly-perceptual belief, as well as the perceptual belief, could be warranted by entitlement.

I shall not further explore the epistemology of such hybrids. I focus on starting points for empirical epistemology: entitlement to perceptual beliefs in my strict sense.

A guiding idea for understanding entitlement—warrant without reason—is that perceptual beliefs can be, and commonly are, epistemically warranted. Perceptual beliefs are not conclusions of propositional inferences, because perceptions are not propositional. Perceptions are not premises in propositional inferences, and are not reasons, because they are not propositional. So perceptual beliefs are not formed by inference from reasons. They are commonly not supported by reasons. They need not be, if believers are to be entitled to them. There is no need for further support than that provided through a reliable, well-functioning, well-grounded competence for forming perceptions and perceptual beliefs from perceptions. That is the primary good route to truth for perceptual belief. That route is commonly sufficient for epistemic warrant. It is often sufficient, in normal situations, to make true beliefs knowledge.

Like justification, entitlement can be *prima facie* or all-things-considered. As with justification, an individual can *have* an entitlement to a psychological state or event, even if it is not operative. For example, an individual can have an entitlement to believe something, but can doubt it—either without warrant or with a justification that defeats the *prima facie* entitlement. As with justification, we can speak of an *individual's* merely having an entitlement or, by contrast, as being (operatively) entitled to a belief or inference.

An individual may *have* both a justification and an entitlement for the same state or event. A state or event can also *be* justified and *have* an *operative* entitlement at the same time. The distinction between justification and entitlement has to do with the nature of types of epistemic warrants. Among the states that one can have a justification or entitlement for is the state of relying on a psychological competence, such as perception, memory, or inference.

The term ‘justification’ is often used for what I call ‘warrant’. This usage is not ideal (see note 17). It need not, however, ruin theory, if the distinction between justification, as I explain the term, and entitlement is kept clear.

Entitlement is not equivalent to what is sometimes called ‘immediate justification’.<sup>28</sup> On this usage, “immediate justification” is warrant that does not derive from

<sup>27</sup> One can certainly test whether a given transition is a propositional inference or an associative transition. Different timings might be involved. Different acquisition histories may be involved. I leave it empirically open whether the non-perceptual concept is directly non-inferentially associated to the *perceptual* body—and shape—attributives, or only to the *conceptual* body—and shape—attributives.

<sup>28</sup> In what immediately follows, I use phrases that involve ‘justification’ in double quotes—scare quotes—to mark others’ use of ‘justification’ in something like the way I use the term ‘warrant’.

other beliefs.<sup>29</sup> Some beliefs are self-evident. A reason to believe them resides in their contents. Understanding their contents suffices to warrant believing them. Examples are I am now engaging in thought (thought occurrently), if something is an impala, it is an impala, and twice two is four. They are “immediately justified”, and also immediately justified, in my usage. Their warrant—their justification—depends for its force on no other propositional attitude or information. Such propositions yield sufficient reason for belief, when believed with minimal understanding. Such a reason is in effect an explanation of the belief’s belief-worthiness. These beliefs are not warranted by entitlement. Perceptual beliefs are not self-evident. All entitlements are “immediately justified”—warranted without reason or propositional inference. Not all “immediately justifications” are warrants by entitlement. The distinction between warrant with and without reason is so fundamental that it should be clearly marked.

Entitlement constitutively depends on both psychological factors and external factors. The external factors are often are not supervenient on the individual’s psychology. The psychological factors are the states or events occurring in a well-functioning competence, though none or these in themselves constitute an entitlement. For example, the perceptual state is not in itself the entitlement. The entitlement is a type of sequence of events or states—from initial registration via perceptual state to perceptual belief—in a reliable, well-functioning competence, that is appropriately connected to a subject matter. A specification of the sequence, its reliability, and its connection to the subject matter would be a reason-explanation of why the belief is epistemically good. To be entitled to the belief, the believer need not be able to produce any such specification.

In the empirical case, the key external factor is reliability in veridical representation of a subject matter. That must be in place if the belief is to count as epistemically warranted. Nothing in the *nature* of the psychological states insures their reliability. Epistemic entitlement does not require that there be warranted meta-beliefs that the perceptions or perceptual beliefs are reliable. In depending on constitutive aspects of warrant, like reliability, that need not be represented, known, or otherwise coded in the mind, entitlement is a partly *externalist* type of warrant.<sup>30</sup> I focus on this matter next.

<sup>29</sup> For examples of this usage, see W. Alston, ‘What is Wrong with Immediate Justification?’ *Synthese* 55 (1983), 73–95; J. Pryor, ‘There is Immediate Justification’ in M. Steup and E. Sosa eds. *Contemporary Debates in Epistemology* (Oxford: Blackwell, 2005); and A. Goldman, ‘Immediate Justification and Process Reliabilism’, in Q. Smith ed., *Epistemology: New Essays* (Oxford: Oxford University Press, 2008). I think use of ‘justification’ in this way is unnatural, though common. See note 17.

<sup>30</sup> Even most *justifications* have an externalist aspect. They constitutively depend on reliability. Deductive inferences from self-justifying truths code this reliability into the nature of the reason-constituting transitions. Perhaps some inductions from self-justifying truths also do so. But most empirical justifications are only *relatively* internalist. The sequence of states/events in the individual’s psychology that constitutes the justification provides an explanatory reason-sketch of why the conclusion is belief-worthy. The sequence of states and events that constitutes a perceptual entitlement provides no such explanatory reason-sketch, because it is not propositional, as explanations must be. In neither empirical justifications nor empirical entitlements are all warranting factors metaphysically necessitated by the nature of the individual’s psychology. I develop these points about perceptual entitlement in Section 2.2.



## 2.2 Empirical Warrant and Reliability

The primary aim of this section is to explain why a modest type of internalism about epistemic warrant is unacceptable. First, I set out the relevant type of internalism. Second, I sketch some history. I explain why, beginning in the early modern period, internalism was a dominant view about epistemic warrant, and how, when its traditional forms of support collapsed in the twentieth century, many philosophers continued to believe it. Third, I criticize some recent misconceptions of epistemic warrant that aim to support internalism. Fourth, I explain why internalism is unacceptable. Finally, I criticize two stock arguments that have been taken in recent times to support internalism. The criticism of the second one explains the specific type of reliability required by epistemic warrant. Accompanying the critical focus of this section is a delineation of central positive aspects of epistemic entitlement for perceptual belief.

The epistemic internalism that I discuss claims that all constitutive aspects of an epistemic warrant are coded in the psychology of the individual who has the warrant. More strictly, I take *epistemic internalism* to claim:

(EI) An individual's having an epistemic warrant, and being epistemically warranted, in having a belief or in engaging in an epistemically relevant transition, supervenes on the non-factive kinds of psychological states, and kinds of relations among them, that are present in the individual's psychology.<sup>31</sup>

That is, if the non-factive kinds of an individual's psychological states and events and their relations to one another are held constant, it is metaphysically impossible for the warrants or specific failures of warrant for those states or relations to vary. Warrant hinges entirely on what psychological states and events are present in the individual with the warrant. It depends not at all on relations between those psychological states and relations, on one hand, and the wider environment in which the individual is embedded, on the other.

Most internalists regard this kind of internalism as modest. Many other internalist positions entail this one. All internalist positions that I know of, including some that do not entail this one, are subject to variants of the criticisms that I make. I discuss (EI) because it is widely taken to be a minimum core of epistemic internalism.

(EI) and most other forms of epistemic internalism are logically consistent with believing that some type of reliable veridicality is constitutively necessary to being warranted. In fact, from Descartes to mid-twentieth century, holding both views was

<sup>31</sup> The limitation to non-factive psychological states is meant to block trivializing internalism by letting such states as knowing count as psychological states. Most formulations of internalist positions are more committal than (EI). Many require that warranting internal psychological states are knowable by reflection from the inside. Many require that one have reflective access, from the inside, to knowledge as to whether a given state or event is epistemically warranted, or reflective access to some other aspect of the meta-explanation of the warrant.

In 'Perceptual Entitlement', *op. cit.*, 504, I characterized epistemic internalism as the view that the epistemic warrant must be conceptually accessible—thinkable—by the individual whose psychological states or events are warranted. I hope to discuss relations between this kind of internalism and the kind characterized by (EI) on another occasion.

standard. If one is to hold both views consistently, one must take one of two approaches. Either one must take reliability to be necessitated by the natures of psychological states and events that are epistemically warranted. Or one must invoke some other metaphysically necessary guarantee of their reliability.

I turn to some history. Beginning with Descartes, many philosophers sought a guarantee that we have knowledge, particularly scientific knowledge.<sup>32</sup> This guarantee was to provide a bulwark against scepticism and a philosophical underpinning for science. In the early modern period, two ways of providing a guarantee dominated. One was to advocate idealism—the view that connection to the physical world just is following epistemically good procedures. Berkeley and Kant held versions of such a view. The other way was to invoke a metaphysically necessary connection between epistemic operations and the world. Descartes, Malebranche, and Leibniz grounded such connections in theology. Both strategies entailed that epistemic warrant is necessitated by an individual's psychological states. As a matter of necessity, one need only form certain types of beliefs and follow certain rules, to be guaranteed to reliably track truth.

Rejection of idealism and theologically grounded metaphysics was basic for Russell and Moore in the early twentieth century. Their rejections turned philosophy decisively away from such failed metaphysical and theological projects. The traditional bases of support for internalism vanished. These philosophers persisted, however, in seeking, against scepticism, a guarantee for empirical knowledge. They built such a guarantee into their accounts of initial empirically relevant psychological states. They took acquaintance with sense-data to ground empirical warrant. They regarded acquaintance as infallibly veridical. They and followers, like Broad and C. I. Lewis, and many positivists, such as Ayer and early Carnap, thought that there were meaning-rules that paved a guaranteed route from sense data to the physical world. The supposed rules caused persistent trouble. Still, the approach dominated epistemology in the first half of the twentieth century. So, like the early moderns, these philosophers tried to show that empirical knowledge, hence empirical warrant, is necessitated by having certain psychological states and following certain rules. In accord with (EI), all features of empirical epistemic warrant, including reliability, were taken to be necessitated by those types of states and rules.

In mid-twentieth century, the ideas that acquaintance with sense data grounds empirically warranted belief, and that one can get from sense data to the physical world via epistemically golden rules, came under devastating attack. Once the strategies of Russell, Moore, and the positivists were abandoned, the historical line of support for (EI) from the early modern period into the twentieth century, was broken. (EI) was ripe for re-evaluation.

At this critical juncture, however, philosophy lacked two things necessary for a successful re-evaluation. It lacked a clear conception of epistemic warrant. And it

<sup>32</sup> As I remarked at the outset, the focus was on a kind of knowledge deluxe, or *scientia*. I am aware that modern uses of 'knowledge' do not match many traditional counterpart uses—uses of '*cognitio*' or '*scientia*', for example. For present purposes, I table these distinctions, although they are historically important. They do not affect the main points that I make here. See note 50. Incidentally, it is possible that for empirical *scientia*, Descartes was an externalist.

lacked an understanding of perception. Realistic reflection on warrant, and indeed knowledge, had lapsed because of the misguided sense-data approaches and because of the battles over positivist theories of meaning. Perception had been ignored because of sense-data theory.

Many post-sense-data philosophers assimilated perceiving to sensing. Sensing was taken not to be representational. So perception was taken to have only causal, not epistemic relevance. Prominent philosophers actually held that nothing in the causal chain before perceptual *belief* is epistemically relevant. Perception disappeared from accounts of empirical warrant and knowledge. Eager to avoid an infallible foundation for knowledge, many required all warranted belief to be backed by reasons. Coherence theories of knowledge and warrant flourished.<sup>33</sup>

Coherence theories were right to emphasize the defeasibility of perceptual beliefs. They seriously underplayed the central role of perceptual belief, let alone perception, in empirical knowledge. They said little about how knowledge and warrant connect to their subject matters. Failure to regard perception and perceptual belief as *prima facie* starting points for empirical warrant led the conception not to mark a connection between empirical warrant and success in representing a belief-independent world.

When opposition to coherence theories gained traction, and perceptual beliefs, broadly understood, regained priority in the structure of empirical warrant, philosophers lacked a thought-through conception of epistemic warrant and a realistic understanding of perception. Relations between environment and psychology that give perceptual states and perceptual beliefs their contents form patterns that knowledge and warrant capitalize on. Such patterns were commonly ignored in the narrow, intra-mural discussions that still dominate post-sense-data epistemology. Many felt free to claim that perceptual belief is intrinsically warranted. Many modest foundationalist approaches continued to accept internalist views of epistemic warrant.<sup>34</sup>

Another factor bolstered this trend. In the post-sense-data period, scepticism's specter continued to warp accounts of epistemic warrant. Many philosophers claimed that warrant must be determinable by resources accessible in one's psychology, on pain of appealing to matters that the sceptic questions. Hasty reflection on sceptical scenarios often led to claims that warrant is independent of psychological-environment relations, and even of any requirement of reliability.

<sup>33</sup> W.V. Quine's epistemology in *Word and Object* (Cambridge, Mass.: MIT Press, 1960), chapter 1 and W. V. Quine and J. Ullian, *The Web of Belief* (New York: Random House, 1970) illustrate these themes. An equally important figure in the development just sketched was W. Sellars, 'Empiricism and the Philosophy of Mind' in *Science, Perception, and Reality* (London, Routledge & Kegan Paul, 1963). Both propounded coherence theories of knowledge. Neither had a clear understanding of perception. Both assimilated perception to sensing. Later, D. Davidson followed Quine and Sellars on these points. See his 'A Coherence Theory of Truth and Knowledge' (1983) in *Subjective, Intersubjective, Objective* (Oxford: Clarendon Press, 2001). More detailed discussion of mid-career Sellars and his followers occurs in 'Perceptual Entitlement', *op. cit.*, section VII.

<sup>34</sup> R. Fumerton, *Meta-epistemology and Skepticism* (Totowa, New Jersey: Rowman and Littlefield, 1995); J. Pryor, 'The Skeptic and the Dogmatist', *Nous* 34 (2000), 517–549; R. Wedgwood, 'Internalism Explained', *Philosophy and Phenomenological Research* 65 (2002), 349–369; J. Cruz and J. Pollock, 'The Chimerical Appeal of Epistemic Externalism', in R. Schantz ed., *The Externalist Challenge* (Berlin: De Gruyter, 2004); S. Berker, 'Epistemic Teleology and the Separateness of Propositions', *op. cit.*

Finally, much of the impetus to believe internalism derived from resistance to “reliabilism”. Reliabilism is widely understood as a view that takes reliability to be sufficient as well as necessary for epistemic warrant.<sup>35</sup> Reliabilism seemed to many to be an extreme form of naturalism that eschews normative, epistemic notions. I do not accept any such reliabilism.

I now criticize some conceptions of epistemic warrant that were conceived to support internalism. These conceptions arose after mid-twentieth century, when the main historical lines of support for internalism collapsed. Traditional lines proposed a metaphysically necessary guarantee for the veridicality of the foundations of empirical knowledge. The approaches that I discuss gave up that project. But they have not been deeply thought through. I think it fair to say that they have not come close to understanding epistemic warrant as a knowledge-conducive epistemic good. Much discussion of epistemic warrant since mid-twentieth century is not guided by any acceptable conception. Confusion reigns.

A view prominent in the heyday of coherence accounts of warrant, which has supporters today, is that all warranted beliefs, including perceptual beliefs, are justified through reasons.<sup>36</sup> Such a view is epistemically internalist only if it takes the warrantedness of a justification’s premises and transitions to depend purely on the individual’s psychology.

The idea that perceptual beliefs are warranted only through an operative argument in the believer’s psychology is wildly implausible. Simple arguments like I saw (or had a perception as of) that red square; so my perceptual belief that that is a red square is (likely to be) true; so that is a red square are justifications. But they do not add to the strength of epistemic warrant for the conclusion beyond what strength entitlement to the perceptual belief already provides. Such arguments are normally neither needed nor relied on. If any more sophisticated arguments were required, few adults would have warranted perceptual beliefs. Young children who have warranted perceptual beliefs and perceptual knowledge (that is body is moving; that surface is red) lack a capacity to think about their perceptions or beliefs as such at all. So they could not think the just-cited argument. Such meta-representational capacities seem required for any perception-based *justifications* of perceptual beliefs. So even simple justifications like the one just cited are probably not available to, much less operative for, all who have warranted perceptual beliefs. More generally, it is wildly implausible that to

<sup>35</sup> D. Armstrong, *Belief, Truth, and Knowledge*, (Cambridge: Cambridge University Press, 1973); A. Goldman, ‘What is Justified Belief?’, in G. Pappas ed., *Justification and Knowledge* (Dordrecht: Reidel, 1979); E. Conee and R. Feldman, ‘The Generality Problem for Reliabilism’, *Philosophical Studies* 89 (1998) 1–29; and A. Goldman, ‘Reliabilism’, *Stanford Encyclopedia of Philosophy*, online. Incidentally, I do not discuss the “generality problem” here. I think it a pseudo-problem that derives from not focusing enough on the actual explanation-grounding kinds of psychological processes that yield given belief states. Most discussion of the issue betrays insufficient use of psychological science in doing epistemology.

<sup>36</sup> Coherence views (see note 33) tended to share with classical foundationalism the view that a warranted individual has a sufficient reason for every warranted belief and hence every piece of knowledge. A coherentist view inspired by Sellars’ ‘Empiricism and the Philosophy of Mind’, *op. cit.* is proposed in L. Bonjour, *The Structure of Empirical Knowledge*, *op. cit.*. Classical foundationalism, unlike most coherence views, maintained that starting points for empirical knowledge, are states that are in effect reasons for themselves. In mid-twentieth century, the idea that there are starting points for empirical knowledge that provide reasons for themselves (or are self-evident or infallible) finally lost credibility.

be warranted and knowledgeable in perceptual beliefs, one must have operative justifications that provide reasons. Most perceptual beliefs do not need argument, or any other justificational background. They are starting points among empirically warranted beliefs and for empirical knowledge.

Another epistemically internalist conception that had a brief run is that epistemic warrant is a matter of being epistemically blameless.<sup>37</sup> This conception is internalist in that it assumes that blameworthiness hinges entirely on states that are in the individual's psychology.

Blamelessness is compatible with systematic malfunction in an individual's psychology. Malfunction undermines warrant. It renders a state incapable of being knowledge. Further, as has often been noticed, benighted superstitions and bad inductive practice can be epistemically blameless. So epistemic blamelessness is not sufficient for having epistemic warrant.

A similar conception is that being epistemically warranted is being epistemically responsible.<sup>38</sup> Epistemic responsibility is a genuine epistemic good. But responsibility, like blame, is only loosely connected to epistemic warrant. Responsibility and blame are relevant only inasmuch as an individual has or has had some control. Some warranted belief formation—including initial formation of perceptual beliefs—is not under an individual's control. Then blame and responsibility are irrelevant to the belief's being warranted. They are not *necessary* conditions. What subsequent epistemic control individuals may exercise on such beliefs does not account for the beliefs' initial prima facie conduciveness to knowledge. Using an unreliable psychological system responsibly also does not *suffice* for epistemic warrant. Indeed, a responsible but unreliable formation of belief would block knowledge, which is an impossible effect of epistemic warrant (see condition (c) in Section 2.1).

In light of the difficulties about control, one could change the view to: being epistemically warranted is either being epistemically responsible or operating in an epistemically permissible way. Even laying aside doubts about whether epistemic permissibility is strong enough for epistemic warrant, the conception still allows cases in which an individual is epistemically warranted, the "warrant" blocks knowledge.

Some internalists take warrants constitutively to provide guidance, recommendation, or advice.<sup>39</sup> They often add that the point of epistemology is to help an individual determine what to believe from the first-person point of view.

Following good guidance is an epistemic good. However, these conceptions are too narrow to yield understanding of knowledge or epistemic warrant. Again, initial formation of perceptual beliefs is commonly not open to following guidance or recommendation.

Often, with crude versions of naturalism and reliabilism in mind, proponents of guidance conceptions insist that epistemology is normative. It is. As I explained in

<sup>37</sup> C. Ginet, *Knowledge Perception, and Memory* (Dordrecht: D. Reidel, 1975); R. M. Chisholm, *Theory of Knowledge*, 2<sup>nd</sup> ed. (Englewood Cliffs, New Jersey: Prentice Hall, 1977).

<sup>38</sup> L. BonJour, 'Externalist Theories of Empirical Knowledge', *Midwest Studies in Philosophy* 5 (1980), 53–73.

<sup>39</sup> J. Pollock and J. Cruz, *Contemporary Theories of Knowledge* (Lanham, Maryland: Rowman and Littlefield, 1999); R. Wedgwood, 'Internalism Explained', *op. cit.*; S. Berker, 'Epistemic Teleology and the Separateness of Propositions', *op. cit.*, and 'The Rejection of Epistemic Consequentialism', *op. cit.*

Section 2.1, epistemic warrant is a normative notion. It fulfills a standard for a type of epistemic success. This standard does not in itself guide. The relevant standard, or norm, is not in itself a prescriptive norm. A common mistake lies in thinking that all epistemic norms, indeed sometimes all norms, are prescriptive norms. Neither guidance nor prescription is essential to warrant or relevant to all types of knowledge.<sup>40</sup> Guidance is relevant to certain types of reasoning. But much epistemic warrant begins below psychological levels at which individuals can be guided.

Guidance conceptions, like blamelessness and responsibility conceptions, take their cue from Descartes' view that epistemology provides rules for the direction of the mind in the interests of gaining knowledge or cognition. (Descartes conceived belief as a product of the will, but proponents of guidance conceptions tend, rightly, not to follow him on this point.) Perceptual-belief formation does not admit of guidance. Moreover, if one's belief-forming equipment is defective or unreliable, one's attempts to follow good epistemic guidance can leave one without warrant or knowledge, even if one's belief is true and not subject to Gettier failures. Indeed, following good epistemic guidance in such cases would still leave one in a condition that blocked having knowledge, which by condition (c) of Section 2.1 is incompatible with being epistemically warranted.

Another internalist conception of epistemic warrant proposes that an individual is epistemically warranted in forming a belief if the belief is likely for that individual to be true, and so, epistemically appropriate for the individual to have the belief.<sup>41</sup> The 'for' and 'so' are not clarified. They are supposed not to require reliability. The natural construal of 'for' is subjectivist. On that construal it simply means that the individual takes it to be reliable, or takes it to be likely true. A madman can meet the condition. So can a believer whose perceptual-belief-forming equipment is pathological. 'For' is perhaps intended to have a more objective status. The conception relies on an unspecified conception of epistemically "appropriate" states or procedures. What makes them appropriate (better, epistemically *good*), and whether any such good is

<sup>40</sup> Epistemology should set out other norms that guide. Discussion of reflective self-monitoring has a place in the epistemology of higher cognition. A common internalist mistake is thinking that all knowledge requires such guidance or in thinking that epistemic warrant must, in general, provide norms of guidance. Epistemic warrant and the basis of empirical knowledge begin below psychological levels at which individuals can follow rules for the direction of their minds.

Arguments for deriving internalism from a guidance-based view of warrant occur in J. Pollock and J. Cruz, *Contemporary Theories of Knowledge*, *op. cit.*, 130–137; and in R. Wedgwood, 'Internalism Explained', *Philosophy and Phenomenological Research* 65 (2001), 349–369, see 365–368. Both arguments state correctly that epistemic evaluation for warrant applies to states and events in an individual's psychology. Wedgwood claims that a transition can be warranted (or "rational", though this term misses entitlement) only if (a) it instantiates a rule that can be directly or non-derivatively followed, (b) it has the force of a "recommendation", and (c) that it is followed because it is warranted (or is rational, or "makes sense"). The argument purports to show that reliability is not a requirement on warrant. The idea is that a recommendation to be reliable cannot be followed directly. I think that the argument provides no reason at all for the conclusion. (b) is not a reasonable condition on warrant. The epistemology of warrant is not confined to recommendation. Moreover, (c) does not accord with how either science or common sense construes formation of perceptual beliefs, many of which are clearly warranted. Nothing in either argument shows that reliability is not a framework-type, necessary condition on a psychological state or process's being warranted.

<sup>41</sup> J. Pryor, 'What is Wrong with Moore's Argument?', *Philosophical Issues* 14 (2004), 349–378.

the good of epistemic warrant are not explained. These are the cruxes of the matter. The relevant epistemic goodness must connect warrant both to truth and to knowledge in ways that I indicated in Section 2.1. Like the previous conceptions, this one does not do so.

Internalist conceptions are often explicated by citing a list of ways of forming epistemically warranted or epistemically appropriate beliefs. Examples are forming perceptions, perceptual beliefs, perceptual memories, inductive inferences, and so on. Citations of such powers, whose exercises in us are typically warranted, do not suffice to illuminate what epistemic warrant is. A general gloss on *wherein* they constitutively conduce to knowledge is lacking. A serious conception of epistemic warrant—or “justification”—is often not developed at all.

Most internalists over the last six decades have lacked any clear, guiding conception of what epistemic warrant is. Post-sense-data conceptions have provided no indication of wherein epistemic warrant is conducive to knowledge. That is the sense in which confusion reigns. In fact, an argument that applies to all the previously cited conceptions is that all allow epistemic warrant or justification to *block* knowledge in certain cases. They all fail to accord with condition (c) in the explication of conduciveness to knowledge that was set out in Section 2.1.

Most of the internalist conceptions of epistemic warrant discussed so far do not directly face the task of defending (EI). This is the task of showing that, and why, an individual’s psychological states and events *necessitate* all features of epistemic warrant. Early modern philosophers confronted the problem directly. They constructed their views about why the mind, by virtue of God’s beneficence or by virtue of the mind-dependent nature of the world, necessarily tracks its subject matters. Many current philosophers are committed, explicitly or in an unarticulated way, to the view that something about the nature of perception or perceptual belief insures that perceptual beliefs are warranted. I do not know of a single substantial or interesting argument in recent philosophy for the position.<sup>42</sup>

<sup>42</sup> For a short time, deplorably and without argument, I claimed that there is such a connection. See T. Burge, ‘Interlocution, Perception, and Memory’, *Philosophical Studies* 86 (1997), pp. 21–47, esp. 31. Shortly thereafter, I rejected the claim. In ‘Perceptual Entitlement’, *op. cit.*, 516–518, 532, I systematically avoid any claim that reliability is constitutive to perception, and I call attention to the contingency of the reliability of perceptual states.

Ralph Wedgwood states that we can know that we are warranted in relying on perception purely by knowing the nature of perception. See ‘Primitively Rational Belief-Forming Processes’, in A. Reisner and A. Steglich-Petersen, eds., *Reasons for Belief* (Cambridge: Cambridge University Press, 2011), section III. The argument for this view seems to me both unclear and flawed in several ways. Either it moves illicitly from the necessity of *some* relevant perceptions’ and perceptual beliefs’ being veridical—a consequence of perceptual anti-individualism—to one’s having a reliable competence. Or it sets a requirement on warrant/rationality that is clearly too weak. The too-weak requirement would be that warrant and rationality bear some constitutive, but not necessarily a reliable, relation to truth. This requirement falls afoul of condition (c) in the explication of conduciveness to knowledge, from Section 2.1. The argument that taking perception at face value is *a priori* warranted has the same flaws. There are other weaknesses in the argument that I will not go into.

Nicolas Silins proposes that perceptual systems, or at least all ranges of perceptual attribution (such as the color range or shape range), are constitutively reliable. See his ‘Explaining Perceptual Entitlement’, *op. cit.*. No argument is offered.

One prominent claim that an intrinsic aspect of an individual's psychology is *sufficient* for warrant states that the phenomenology of perception and perceptual beliefs is the source of perceptual warrant.<sup>43</sup> A given type of perceptual phenomenology is supposed to make a perceptual state or perceptual belief with that phenomenology warranted. Perhaps some advocates think that phenomenology necessitates reliability. Many deny that reliability is required for epistemic warrant. They claim that perceptual phenomenology (constitutively) suffices for default epistemic warrant for perceptual belief. In any case, phenomenology itself does not necessitate reliability and does not constitutively (hence necessitate) conduciveness to knowledge, as epistemically warranted states must. Indeed, if a given phenomenology were not reliable, it would block knowledge, again contrary to condition (c) from Section 2.1.

The only support given for this sufficiency view is that it is "intuitive". It is intuitive that our conscious perceptual beliefs are warranted. It is intuitive that we commonly use perceptual phenomenology in our reliance on perceptual beliefs. We are right to do so. It is also intuitive that our perceptual beliefs and inductive inferences from perceptual beliefs are mostly *prima facie* warranted, without need for supporting argument or background beliefs or presumptions.

None of these points implies that phenomenology insures epistemic warrant, or that conscious perceptual beliefs are *by nature* epistemically *prima facie* warranted. That is the issue. These internalists claim that necessarily epistemic warrant depends on nothing beyond the individual's psychology. Here the relevant aspect of the

[Note 42 continued:]

I discuss reasons to deny such proposals below. Cf. also *Origins of Objectivity*, *op. cit.*, chapters 8 and 9 and 'Perception: Where Mind Begins', *Philosophy* 89 (2014), 385–403.

Ernest Sosa, in *Reflective Knowledge* (Oxford: Clarendon Press, 2009), 236, claims that when we think of a disposition as a competence, we require it to be reliable. He has a perceptual-belief-forming competence specifically in mind. See also his *A Virtue Epistemology* (Oxford: Clarendon Press, 2007), 106. This claim is mistaken. There is no oddity in thinking that someone has a competence in shooting a basketball—or in piano playing, or in archery—but that the competence is erratic and unreliable, and not just because it is interfered with by external factors. An ability is not a probability. More specifically, there is nothing impossible about a perceptual or a perceptual-belief-forming competence that is not reliably veridical in its natural habitat. Empirical science proceeds on this assumption. It investigates whether, to what degree and in what circumstances, perceptual competencies are reliable.

<sup>43</sup> Representative proponents are J. Pryor, 'The Skeptic and the Dogmatist', *op. cit.*; M. Huemer, *Skepticism and the Veil of Perception* (Lanham, Md.: Rowman & Littlefield, 2001) and 'Phenomenal Conservatism and the Internalist Intuition', *American Philosophical Quarterly* 43 (2006), 147–158; M. Steup, 'Internalist Reliabilism', *Philosophical Issues* 14 (2004): 403–425; and I. Dickie, 'Visual Attention Fixes Demonstrative Reference by Eliminating Referential Luck' in C. Mole, D. Smithies, and W. Wu eds., *Attention: Psychological and Philosophical Essays* (Oxford: Oxford University Press, 2011). Dickie declares that intentionality is 'essentially a matter of subjective consciousness', citing as opponents only naturalistic philosophers who envisage reduction of intentionality to non-intentional, non-mentalistic phenomena (299–300, 315). Such declarations are incompatible with what is known from the science of perception or perceptual-belief formation. Science freely postulates unconscious representational states (states with "intentionality"). It does not appeal to consciousness as the basis for explaining successful perceptual reference or perceptual attribution, even in cases of states that are conscious.

There is an aberrant version of the view that phenomenology is sufficient for warrant that *identifies* phenomenology with environmental objects and attributes. I do not take this version seriously. See note 127.



psychology is supposed to be the conscious phenomenology of the individual's perceptual states or perceptual beliefs. An account of wherein the phenomenology bears a *constitutive*, or at least *necessitating*, relation to epistemic warrant—hence to conduciveness to knowledge—is completely lacking.

The claim that phenomenology is what makes a perceptual state or perceptual belief epistemically warranted is hocus pocus. No proponent, to my knowledge, has attempted a serious explanation of how any specific empirical phenomenology, or consciousness, could in itself yield epistemic warrant. No one has explained wherein a phenomenology is constitutively conducive to knowledge. How phenomenology could constitutively determine these matters is mysterious.

Perception is fundamentally a functional matter. A tendency to realize its representational function in representing the world correctly is its key contribution to epistemic warrant for perceptual belief. Although consciousness *actually* enhances veridicality and reliability, nothing about phenomenology is *constitutive* to this representational function.

I have criticized the view that conscious phenomenal experience *suffices* for having empirical epistemic warrant. I shall return to that view. As an interlude, I discuss a claim that conscious phenomenal experience is *necessary* for having epistemic warrant. John Campbell argues that conscious phenomenal experience—indeed conscious focal attention—is *necessary* for having demonstrative reference to objects in thought, hence to having “justified” perceptual belief. He thinks that to “grasp” or “understand” a demonstrative in thought, one must consciously attend to the object that the demonstrative picks out in perceptual experience.<sup>44</sup> I believe this view to be mistaken. It is certainly not successfully argued.<sup>45</sup>

<sup>44</sup> The view that focal attention is required for having demonstrative-governed perception-based beliefs was inspired by scientific work by Anne Treisman. In other work, *Perception: First Form of Mind*, forthcoming, I explain how the view has come to be undermined by psychological science itself.

<sup>45</sup> J. Campbell, *Reference and Consciousness* (Oxford: Clarendon Press, 2002), 7–9. D. Smithies, ‘What is the Role of Consciousness in Demonstrative Thought?’, *The Journal of Philosophy* 108 (2011), 5–34, especially 33ff., endorses Campbell’s argument and rests further arguments on it. I. Dickie, ‘Visual Attention Fixes Demonstrative Reference by Eliminating Referential Luck’, *op. cit.* also is influenced by Campbell’s claims. These authors assert that demonstrative reference in thought, at least when backed by categorial attributives, can occur only consciously, and that no unconscious demonstrative-involving belief can be warranted. I do not accept these assertions. Unconscious perception can and does attribute categorial attributes, like size and shape—for example in blindsight and in certain neuro-typical cases of human perception. See the citations in this note regarding blindsight. Unconscious perceptual beliefs sometimes are formed from such perceptions. Unconsciousness does not disqualify categorial perceptual beliefs from involving demonstrative reference—from picking out entities in perception via contextual causal relation to them—or from being warranted, although actual blindsighters may form beliefs that are too weakly supported to count as warranted.

Ian Philips doubts even the view that there *is* unconscious perception. Standard perceptual psychology postulates many cases of unconscious perception. Philips criticizes that standard view, and my and Ned Block’s acceptance of it, in ‘Unconscious Perception Reconsidered’, *Analytic Philosophy* 59 (2018), 471–514. No arguments are presented to show that consciousness is *constitutively necessary* to perception. The criticisms are empirical and case by case. Philips uses my characterization of perception as objective sensory representation by the individual as the framework for his criticisms (*Origins of Objectivity, op. cit.*, 368–416).

Philips’s main claim is that there is no unconscious perception *by individuals*. Philips invokes my notion by individuals. On my understanding, what it is for a perception to be *by an individual* is for it to be functionally integrated with exercises of whole individual functions. Such exercises can be active, such as turning one’s

[Note 45 continues below next page]

[Note 45 Continued:]

head, or non-active, such as automatic freezing in response to danger. Philips's holds that there is no evidence for unconscious perception that is not functionally isolated from exercises of whole individual functions.

However, there is considerable evidence that unconscious perception guides action—say, in blindsight—and guides allocation of attention—whether active or “grabbed” attention. Philips assumes, I think mistakenly, that only active, guided attention is attributable to an individual. I think that in all relevant cases, active or “grabbed”, attention is attributable to an individual. There is also evidence that unconscious perception functionally interacts in systematic ways with later conscious perception by the individual. Such unconscious perception is thus closely integrated with whole-individual functions—with further perception and whole-individual responses to perception, including actions. I find Philips' applications of my notion “by an individual” are often quite foreign.

Philips's *focus* on whether a phenomenon is perception by *an individual* is, I think, itself misconceived. Ned Block, who disagrees with most of Philips's other contentions, shares this focus. See N. Block, 'The Anna Karenina Principle and Skepticism about Unconscious Perception', *Philosophy and Phenomenological Research* 93 (2016), 452–459; and I. Philips and N. Block 'Unconscious Seeing'—a debate in B. Nanay (ed.) *Current Controversies in Philosophy of Perception*, (London: Routledge, 2016). Both authors take whether a visual state is attributable to an individual to be crucial to whether unconscious and conscious sensory representations are sub-species of a basic or important psychological natural kind. I disagree. I explicitly avoided taking being “by an individual” to be necessary to perception (*Origins of Objectivity*, *op. cit.*, 369). 'By the individual' figures in my explication of perception because perception is paradigmatically by individuals. The phrase is a generic. It does not set a necessary condition. Perceptual constancies, conscious or not, and also by the individual or not, are necessary and sufficient for perception (*Origins of Objectivity*, *op. cit.*, 413). Such constancies mark a basic, important psychological natural kind. The basic natural kind is *perception!* Conscious and unconscious species, whether by the individual or not, are sub-kinds.

Some blindsight clearly guides individual action. So some blindsight is clearly by individuals. In supporting criticisms of what I had written on blindsight, Philips's reporting of science is very unreliable. He writes, 'we lack evidence that blindsight involves genuine perception [perceptual constancies] as opposed to mere sensory registration.' He cites two studies of blindsight that suggest that specific types of visual detection—sensing motion not tracked in three-dimensional space, and sensing certain two-dimensional shapes—do not involve perceptual constancies. I think that these cases do not suggest, even *prima facie*, the presence of perceptual constancies. Philips generalizes from these cases as if they were the only relevant ones. He passes over, unmentioned, the very cases that I cited (*Origins of Objectivity*, *op. cit.*, 374–375). Blindsight subjects' grasping 3-dimensional objects, via smooth motion and with fitted aperture of grasp, constitute the sort of evidence that supports attribution of perceptual constancies. A large scientific literature claims or strongly suggests perceptual constancies in blindsight. See 'J. Dankert and Y. Rossetti, 'Blindsight in Action: What Can the Different Sub-Types of Blindsight Tell Us about the Control of Visually Guided Actions?', *Neuroscience & Biobehavioral Reviews* 29 (2005), 1035–1046; M. Perenin, and Y. Rossetti, 'Grasping Without Form Discrimination in a Hemianopic Field', *Neuroreport* 7 (1996), 793–7. M. Christensen, L. Kristiansen, J. Rowe, and J. Nielsen, 'Action-blindsight in Healthy Subjects after Transcranial Magnetic Stimulation' *PNAS* 105 (2008), 1353–1357; R. Whitwell, C. Streimer, D. Nicole, and M. Goodale, 'Grasping the Non-conscious: Preserved Grip Scaling to Unseen Objects for Immediate but not Delayed Grasping Following a Unilateral Lesion to Primary Visual Cortex', *Vision Research* 51 (2011), 908–924; R. Whitwell, A. Milner, C. Cavina-Pratesi, M. Barat, and M. Goodale, 'Patient DF's Visual Brain in Action', *Vision Research* 110 (2015), 265–276.

There are direct arguments for perceptual constancies in blindsight: I. Sperandio and P. Chouinard, 'The Mechanisms of Size Constancy' *Multisensory Research* 28 (2015), 253–283; R. Whitwell, I. Sperandio, G. Buckingham, P. Chouinard, and M. Goodale, 'Evidence for a Functional and Anatomical Dissociation in the Use of Size Constancy for Perceptual Report and Goal-directed Grasping' (abstract), *Journal of Vision* 15 (2015), 187.

Philips gives back-up arguments that purport to show that none of the evidence used in science to support taking blindsight subjects to be relevantly unconscious really does so. I find the arguments unpersuasive. They join a long line of unsuccessful, rear-guard attempts to show or suggest that blindsight patients always perceive consciously. I think that the science will not be moved by Philips's attempts along this line. Moreover, as Block indicates, there are many cases of unconscious perception by individuals who are *neuro-typical*.

Campbell's argument goes as follows:

You might acknowledge that ordinarily we would use visual information to interpret the demonstrative, but question whether it has to be conscious. The idea of visual information that is not conscious is made vivid by cases of blindsight. A blindsight patient is one who . . . has no awareness of objects in one half of his visual field. Nonetheless, when forced to guess about what is in the blind field, he may be reliably correct about, for example, the orientation, direction, and sort of object in the blindfield. . . . Couldn't this subject use such visual information to achieve an understanding of a visual demonstrative? . . . we could try supposing as a thought experiment that we have a subject who has all the relevant visual information but is not yet conscious of the object in the blind field ( . . . 'super-blindsight'). [Campbell goes on to claim that super-blindsight is just guessing.]<sup>46</sup>

. . . consider an ordinary case in which you and I are sitting at a dinner table with a large number of people around and you make a remark to me about 'that woman' . . . I cannot consciously single out the person you mean. All I get consciously is a sea of faces. But now we add some of what the blindseer has. You refuse to give me any further clues as to which person you mean, but you say, 'Try to point to the woman I mean'. At first I protest that I can't do that . . . but I do try to point, and to my surprise you say I'm pointing right at the person you mean. Suppose now that my conscious experience remains a sea of faces, but extend the reach of my reliable guessing so that it encompasses everything the blindseer can do. So I can make reliable guesses about what the person is eating, wearing, and so on, as well as reaching and pointing appropriately. But so long as my conscious experience remains a sea of faces, there is an ordinary sense in which I do not know what you mean. The problem here does not have to do with whether I am reliable. We can suppose that I am quite reliable in my guesses and we establish this over a series of such cases. The point is rather that I do not know who you mean until I finally look at where my finger is pointing or look to see who is wearing the clothes I described in my guesses. It is only when I . . . single out the woman in my experience of the room, when it ceases to be a sea of faces, and in my experience I focus on the person, that I would ordinarily be said to know who was being referred to. So it does seem to be compelling to common sense that conscious attention to the object is needed for an understanding of the demonstrative.<sup>47</sup>

<sup>46</sup> J. Campbell, *Reference and Consciousness*, op. cit., 144.

<sup>47</sup> J. Campbell, *Reference and Consciousness*, op. cit., 7–9. Campbell introduces this passage in misleading ways. He writes, 'The issue is whether the blindseer has the very same way of interpreting the demonstrative as the ordinary subject has. That is, the question is whether for the ordinary subject consciousness of the object is not completely idle in an understanding of the demonstrative' (8). These formulations are not equivalent to one another. And neither formulation is equivalent to the issue that the argument addresses—whether consciousness is constitutively necessary to having demonstrative thought. Anyone should agree that an unconscious "interpretation" (or mode of presentation or way of perceiving) differs from a conscious one. Moreover, the issue is not whether consciousness is 'completely idle' in a conscious individual's understanding the demonstrative. The issue is whether it is constitutively necessary for demonstrative, *de re* thought. Consciousness quite obviously plays a role in actual competencies with conscious *de re* perceptual thought. It is empirically obvious that consciousness sharpens and enriches perceptual information in many uses. Its psychological role is hardly idle. Conflation between being

The argument fails in two ways. One centers on characterizing the super-blindseer as guessing. The super-blindseers are supposed to have reliably veridical perception and to make judgments with confidence. *Actual* blindseers *take* themselves to be guessing. But both they and super-blindseers have a real competence. A lot is known about the competence, and about the perceptual constancies involved. Given *super-blindseers'* competence at perceiving bodies and shapes as such, their reliability, and their confidence, they are not guessing.

Campbell insists that even 'fast, accurate, and reliable' super-blindseeing 'is still said to be 'only guessing'.' (See note 46.) But where is it "said", outside this insistence, that a *super-blindseer* would be guessing? Dictionaries explicate guessing as taking something to be correct without sufficient information to be confident about it. Super-blindseers' competent, confident, information-based judgments fail this condition. Campbell's argument cannot draw on shared common sense to say that a competent, confident, reliably veridical super-blindseer would be guessing.

I and perceptual psychology (see citations note 45) take the *ordinary* blindseer to be perceptually picking out an object and attributing properties like shape and location to it. The ordinary blindseer has unconscious demonstrative thoughts as of categorial properties formed from unconscious perceptions of, and as of, the object. I think that the super-blindseer's thoughts—given their high reliability, confidence, and naturalness—would be knowledge. No attributions of guessing to the super-blindseer have any force.

A second failure of the argument lies in its comparison of super-blindsight with attempts to understand the linguistic demonstrative references of another individual who gives no cues as to which face is meant or demonstratively referred to in thought. There is no mechanism or competence that would, in the absence of such cues, enable one to track which face the other individual intends. We do not read others' minds except through perceived cues. So there is no way for the hearer in the sea-of-faces thought experiment to understand and competently track the speaker's demonstrative usage, or to replicate it in thought. Here the situation differs crucially from blindsight and super-blindsight. Blindseers and super-blindseers have a known competence for tracking objects that they unconsciously perceive.

The goal of Campbell's argument is to show that conscious attention is necessary for the blindseers' use of demonstratives in thought. To be successful, the argument must exhibit a case in which demonstrative thought is absent and the only feature that can account for its absence is absence of conscious attention. In the sea-of-faces thought experiment, conscious attention is absent. Its absence is supposed to explain why demonstrative thought is (supposed to be) absent. Also absent is a mechanism or competence for tracking and replicating the other's demonstrative reference. This latter absence suffices to explain the hearer's failure to think demonstratively about what the speaker means. The argument does not show that consciousness or a conscious mechanism is what explains the failure of the hearer to secure

constitutive to *de re* reference and being non-idle in normal, conscious *de re* reference figures in other philosophical discussion of these issues.

demonstrative reference in the sea-of-faces case. Inferring the conclusion from the thought experiment is a *non-sequitur*.<sup>48</sup>

Campbell frames the discussion as if it were analogous to understanding another's speech. But relevant understanding for perceptual beliefs is different. It is a competence to rely on one's own perception in contextually picking out objects in thought. The super-blindseer has relevant understanding of such thoughts—competence in thinking them. Campbell has not given the slightest reason to think that the super-blindseer lacks relevant competence grounded in unconscious perception. Since *ordinary* blindsight is a defect, the individual is disinclined to rely on it. Ordinary blindsight is not strongly reliable. Hesitancy to form beliefs on the basis of ordinary blindsight and lack of strong reliability prevent it from yielding knowledge. The hypothetical super-blindseer, whose blindsight is natural, confident, and reliable, understands his or her thoughts in the only way that is relevant. Relevant understanding of demonstratives in thought consists in cognitive competence to think thoughts that pick out objects via perception. The demonstrative-like referential applications in the thoughts and perceptions pick out objects contextually, in ways that are embedded in perceptual constancies.<sup>49</sup>

In sum, the argument that consciousness is constitutively necessary to having warranted, demonstrative-governed, perception-based thoughts clearly fails. Although actual blindseers lack knowledge, because their perceptions and perceptual beliefs are not reliable enough and because they lack confidence in their (forced) judgments, hypothetical super-blindseers' perceptual judgments have not been shown to lack warrant or knowledge.

I have criticized both the view that consciousness in perception is sufficient and the view that consciousness in perception is necessary for epistemic warrant in perceptual belief. The claim that conscious phenomenology suffices for epistemic warrant is by far the more pernicious. The hopelessness of the position follows from the fact that not even the natures of perception and perceptual belief, or any given perception or perceptual belief—let alone consciousness in them—constitutively insures epistemic warrant. I shall discuss this point as part of a general explanation of why internalism is unacceptable.

<sup>48</sup> Here the weakness in Campbell's argument is similar to a weakness in Laurence Bonjour's argument from clairvoyance, which I discuss later in this section.

<sup>49</sup> Some of Campbell's argument hinges on the entirely armchair claim that a blindseer *only* perceives functional properties, "affordances", *Ibid*, 138–145. It is an open empirical question whether *any* functional attributives are formed in the visual system, as distinguished from the actional system. I assume that some are. Campbell's armchair claim that blindsight involves perception *only* of affordances is, however, flatly incompatible with science. Shape constancies and categorizations of shape and shaped-body are attributed in the science to ordinary blindseers. See R. Whitwell, C. Streimer, D. Nicole, and M. Goodale, 'Grasping the Non-conscious: Preserved Grip Scaling to Unseen Objects for Immediate but not Delayed Grasping following a Unilateral Lesion to Primary Visual Cortex', *op. cit.*; I. Sperandio and P. Chouinard, 'The Mechanisms of Size Constancy', *op. cit.*; R. Whitwell, I. Sperandio, G. Buckingham, P. Chouinard, and M. Goodale, 'Evidence for a Functional and Anatomical Dissociation in the Use of Size Constancy for Perceptual Report and Goal-directed Grasping', *op. cit.*. Moreover, Campbell's claim is incompatible with the *main shape* of modern vision science (see note 45). Modern (post-1970) mainstream vision science never takes perceptual attribution of functional properties to be independent of perceptual attribution of categorical properties. One perceives something as prey only by perceiving it as having such attributes as shape, size, and body.

The explanation is what is most important. I first state an argument that it frames. Then I explain and motivate its component steps:

- (1) Many individuals are epistemically warranted in, entitled to, many of their perceptual beliefs.
- (2) An individual's perceptual belief's being produced or sustained in a way that reliably engenders true beliefs of that belief's content-type is constitutively, hence metaphysically, necessary to the belief's being epistemically warranted. (Reliability is to be assessed in the perceptual and belief systems' normal environment, in a sense of 'normal' to be discussed later in the section.)
- (3) For a psychological system of states that actually has a warranted perceptual belief that has actually been produced or sustained in a way that reliably engenders true beliefs of that belief's content-type, it is not metaphysically necessary that a belief of the same content-type that is embedded in the same configuration of psychological state types, and that is produced or sustained in the same way, is reliably true. (Nor therefore is the perceptual belief reliably true by its nature or content.)

Hence

- (4) It is metaphysically possible for the whole configuration of psychological state types that embeds an epistemically warranted perceptual belief to remain the same, while a perceptual belief of the same type is (would be) epistemically *unwarranted*.

(4) entails that the epistemic warrant, the entitlement, to a perceptual belief does not supervene on the system of psychological states in which it is embedded, contrary to (EI).

I have already supported (1). Rejection of (1) in the interest of requiring propositional buttresses for perceptual beliefs is deeply implausible. Young children have warranted, knowledgeable, perceptual beliefs, such as that round body is red, or that is a face. They lack reasons and background beliefs that function to support, protect, or buttress their perceptual beliefs. Adult epistemic practice presumes no need for propositional support or protection for perceptual beliefs. Denial of these points often evinces bad theory's swamping good, common sense judgment. Or denial evinces a focus on a special type of knowledge or epistemic warrant that not all individuals with knowledgeable perceptual beliefs have. I discuss these matters further at the end of this section and in Sections 2.3–2.5. Although I think (1) clearly true, a more complex argument could be refashioned from steps (2) and (3) alone.

In this section, I focus on (2) and (3). I believe that they are the central claims at issue. I think that most epistemic internalists would be disposed to dispute either (2) or (3), or both. I explain why such dispositions should be reconsidered.

I begin with the claim in (2):

An individual's perceptual belief's being produced or sustained in a way that reliably engenders a true perceptual belief of that belief's content-type is constitutively, hence metaphysically, necessary to the belief's being epistemically warranted.

I argued for this claim in Section 2.1. I elaborate the argument here. Later in this section, I develop the *relevant* type of reliability required for epistemic warrant.

Epistemic warrant's being *epistemically* good consists in its conducing to knowledge—in its marking a belief or inference as being produced or maintained through a competence that tends to be, or to contribute to having, knowledge. If a belief were produced or sustained by a competence that was in every relevant respect unreliable, the belief would not be formed or maintained in a way that *could* be knowledge, much less tend toward having knowledge. Its being formed in that way would in fact *prevent* the belief from being knowledge. Epistemic warrant would in those cases mark a decisive hindrance to having knowledge. Such a result—a violation of condition (c) on conduciveness to knowledge from Section 2.1—prevents any property that could possibly block knowledge from being *epistemic* warrant. Constitutively, epistemic warrant marks a state or process as being situated in a pattern of psychological competencies that conduces to having knowledge. It is impossible that one lacks knowledge *by* being epistemically warranted. One cannot ignore the connection between the concepts epistemic and knowledge. So all epistemic warrant constitutively requires that the relevant competence for producing or sustaining a warranted belief be relevantly reliable in yielding true belief. An analogous argument applies for epistemically warranted inferences, or other transitions, to a belief.

In many cases, an epistemic warrant certifies true beliefs *prima facie* as knowledge. Among these cases are epistemically basic beliefs—including perceptual beliefs. In such cases, the basic or canonical warrant (perceptual entitlement, for perceptual beliefs) certifies true beliefs as knowledge, in the absence of certain undermining conditions. I conjecture that these conditions are an absence of counter-warrant, an absence of internal psychological incoherence, and an absence of Gettier conditions. A true belief could never be certified as knowledge by epistemic warrant, even if it met such absence conditions, if it were produced or maintained by a competence that was relevantly unreliable. So basic epistemic warrants for basic beliefs could not certify those beliefs, even *prima facie*, as knowledge if epistemic warrant did not require that the competence for producing and maintaining the beliefs be relevantly reliable.

Many internalists cheerfully deny that any sort of reliability is a condition on a belief's or inference's being epistemically warranted. Some offer substitute conditions—that the individual *not believe* that a warranted state is produced in an *unreliable* way, or that the individual have a warranted *belief* or *presumption* that the state is produced in a reliable way.

Denials that reliability is a constitutive aspect of epistemic warrant derive from confused failure to recognize that epistemic warrant is an *epistemic* good in that it constitutively conduces to knowledge. The epistemic good, epistemic warrant, cannot in any situation block or prevent knowledge of the warranted belief by the believer. A belief's being produced in a way that does not reliably engender true belief blocks knowing it. This consequence is *a priori* unacceptable. Such a belief could not be epistemically warranted.

Such denials are historically aberrant. No philosopher in the tradition would have taken seriously a denial that his analog of warrant or justification must be reliable in

marking true belief.<sup>50</sup> Each would have taken it as basic that a good procedure for gaining knowledge, or *cognitio*, or *scientia*, or *Erkenntnis* must be reliable. Each would have recognized that a way of producing or sustaining a belief or inference that is relevantly unreliable *could not* mark a true belief as knowledge, and *could not* conduce to having knowledge. Epistemic warrant constitutively marks a true belief as conducing to knowledge. So an epistemically warranted procedure for gaining knowledge must be relevantly reliable, in producing truth. Denial that a warranted belief must be produced or sustained in a reliably veridical way betrays conceptual confusion about what epistemic warrant is.

It would be seriously wrongheaded to respond in the following way.

The early-modern internalist tradition of guaranteeing that analogs of epistemic warrant are conducive to knowledge and reliable in yielding truth was an optional aspect of the tradition. The commitments to idealism and metaphysical schemes that claimed that certain good ways of empirical thinking must necessarily yield knowledge and understanding were optional. What mattered to the early modern philosophers was their internalism and their articulating epistemically responsible procedures. So the right understanding of their analogs of epistemic warrant is an internalist conception that centers on responsibility.

Making such a position explicit should bring out how misguided, and unhistorical it is. One needs to know only a little of the history of early modern philosophy to recognize that this gloss is uninformed. The early moderns aimed at understanding the new science—understanding how thinking scientifically contributed to arriving securely at an understanding of the structure of the world. Understanding our scientific tracking of the world was Descartes' primary objective. With the exception of Hume, the early moderns overestimated how much of this tracking could be understood through reflection alone. Many of them felt more confident than we can be in appealing to theological or other metaphysical strategies to explain such tracking. They thought that by articulating procedures that fitted with their metaphysical schemes, they were showing how minds were guaranteed to track the world, at least in its larger patterns. Internalism was a consequence. But the primary objective in exploring analogs of epistemic warrant was to understand what psychological states and procedures conduced to their analogs of knowledge.

<sup>50</sup> As I mentioned earlier (note 32), the philosophical tradition rarely focused on knowledge. It focused on *scientia* or (for Kant) *Erkenntnis* ("cognition"). The former is very roughly analogous to scientific knowledge. Kant distinguished *Erkenntnis* ("cognition") from *Wissen* (knowledge). For him *Erkenntnis* is a complex notion. In the form that interested Kant most, it is also *very* roughly analogous to a high-level type of knowledge, in that it requires an ability to *prove* objective validity (*Critique of Pure Reason*, Bxxviii). I do not know whether before Kant, there was use of terms like 'warrant' or 'justification'. Kant has a close, counterpart notion for justification—proof of objective validity. Before Kant in the early modern period, the analog of warrant or justification was good procedure for obtaining *scientia*. Descartes certainly focused on such procedures. He and Kant regarded their analogs of warrant as reliable in producing truth and in producing their analogs of knowledge. A twentieth century recovery of the insight that reliability is a condition on warrant, and on any good procedure for yielding knowledge or *scientia*, occurs in Alvin Goldman, 'What is Justified Belief?', *op. cit.*. Without endorsing all of Goldman's positive view, I believe that his emphasis on reliability as a necessary condition on being warranted ("justified") has been important and salutary.



The ethics of epistemically responsible practice was set out by Descartes to serve an understanding of science's tracking the world. Indeed, except for Descartes, none of the early moderns featured notions of epistemic responsibility in their accounts of (analogs of) knowledge. The central focus, even in Descartes, was how thought's, especially scientific thought's, tracking the world could be explained and understood.

Modern internalists who think that an individual or state can be epistemically warranted, but be produced or sustained in a way that is unreliable in producing true beliefs, propose a notion of epistemic warrant that could consistently be a hindrance to knowledge. The early moderns would have and should have regarded any such notion as *apriori* unacceptable as a notion of epistemic warrant.

It is certainly true that there are coherent internalist notions of being responsible, or of being blameless, or of using one's psychology as best one can, in serving knowledge and true belief. These notions do have uses in understanding norms for some epistemic practices. They are not central to understanding the factors that make one's psychological states good in knowing the world. As I have noted, the blame/responsibility notions do not even apply naturally to the most basic way of forming knowledgeable beliefs—entitled production of perceptual beliefs. Initial formation of perceptual beliefs, and epistemic warrants for them, begin below the level of norms of blame, responsibility, or guidance.

How does the deep connection between epistemic warrant and knowledge bear on the controversies over defining knowledge? Does such a connection block attempts to *analyze* the concept of knowledge in terms of truth, epistemic warrant, and other concepts? Does it support the idea that epistemic warrant is to be defined in terms of knowledge? No, to both questions.

I do not purport to analyze or define concepts as complexes of other concepts. I doubt common understandings of these enterprises. I think that most important philosophical concepts are not definable. They are among the conceptual starting points for understanding the world. However, the claims set out here are not incompatible with defining knowledge partly in terms of epistemic warrant. The connections between epistemic warrant and knowledge that I have articulated are certainly not obviously definitional or analytic connections. They are also not incompatible with an attempt to analyze epistemic warrant in terms of knowledge. I do not propose definitions of either knowledge or epistemic warrant. Specifically, I do not hold that the concept knowledge is part of the definition of the concept epistemic warrant, or that knowledge is otherwise prior to epistemic warrant in some order of understanding.<sup>51</sup> I tend to think that the concepts of knowledge and

<sup>51</sup> Thus my view is not intended as a form of “knowledge first” approach to epistemology. See T. Williamson, *Knowledge and Its Limits* (Oxford: Oxford University Press, 2001). On such approaches, epistemic warrant and the psychological competencies, states, operations, and structures that fall under epistemic norms tend to be ignored. Such approaches tend to center on counterfactual, “non-luck” accounts of relations between beliefs and the world in attempting to illumine the concept knowledge. While these accounts can be valuable in understanding counterfactuals that underlie conduciveness to truth, they provide little understanding of the psychological bases for knowledge—how psychological competencies, states, operations, and structures contribute to gaining knowledge. Nor do they tend to provide insight into epistemic norms that govern the relevant psychological phenomena. Epistemology

epistemic warrant are primitive, but apriori reciprocally inter-dependent. I am doubtful about definitions of important concepts (or analysis of the meaning of terms) that play such rich roles in our understanding of the world.

I do believe that it is apriori that knowledge requires epistemic warrant and true belief. I also believe that it is apriori that epistemic warrant marks conduciveness to knowledge and reliability regarding truth. Most important apriori truths are synthetic in the sense of not being purely products of analysis of concepts, on an analogy with analysis of meaning. (I think that *all* truths are synthetic in the sense that their truth depends on being made true by a subject matter.)

What is clear is that no property that does not conduce to knowledge and reliably signal truth could be the epistemic good that notions of epistemic warrant, entitlement, and justification in fact mark. No property that would prevent knowledge could be epistemic warrant. Epistemic warrant constitutively marks psychological states as deriving from competencies or procedures that yield, sustain, or contribute to knowledge.

I turn to step (3) in the argument against (EI):

For a psychological system of states that actually has a warranted perceptual belief that has actually been produced or sustained in a way that reliably engenders true perceptual beliefs with that belief's content-type, it is not metaphysically necessary that a belief of the same content-type that is embedded in the same configuration of psychological state types, and that is produced or sustained in the same way, is reliably true. (Nor therefore is the perceptual belief reliably true by its nature or content.)

Reliability is coded into the very *nature* of a few types of psychological states and events. A belief in a logical or mathematical truth is reliably true by its nature. Of course, belief in such a truth on a whim does not derive from competence that is connected to the belief's subject matter in a reliable way. Accepting a mathematical truth on a whim would be reliably true by the nature of the belief, but not through the nature of a competence that reliably connects to the belief's being true. So a belief in a logical or mathematical truth on a whim is not epistemically warranted. However, understanding the representational content of such a belief can constitute a good, reliable route to true belief. For relatively simple logical and mathematical truths, their content is self-evident: understanding their content suffices to warrant belief in them.<sup>52</sup> So believing on the basis of understanding the content makes one's belief

must illumine both matters. Although I share Williamson's doubts about twentieth-century attempts to define the concept of knowledge and his rejection of epistemic internalism, I do not find his arguments for the "knowledge first" view clear or persuasive. Beyond these remarks, I will not discuss that work or attempts by others to develop it. I mainly caution against taking my point that the concept of epistemic warrant bears apriori relations to the concept of knowledge to be associated with the slogan, "knowledge first". Although there is certainly value in approaches motivated by this slogan, I regard it, and them, as uncongenial in their way of pursuing epistemology.

<sup>52</sup> I support this view in *Truth, Thought, Reason: Essays on Gottlob Frege* (Oxford: Oxford University Press, 2007), 65–72, 301–316; and 'Logic and Analyticity', *Grazer Philosophische Studien* 66 (2003), 199–249.

formation not only reliably true. It makes it reliably true by a route that is reliably connected to the subject matter that makes the belief true. Reliability is coded into accepting self-evident necessary truths on the basis of understanding them. Similar points apply to deductively valid inferences. An inferential transition made via a competence individuated by a deductive inference rule is warranted by its nature. Epistemic warrant is necessitated by these types of psychological states and events.

Empirically warranted psychological states and events are different. There is no intrinsic guarantee of their reliability. In the first instance, reliability is a property of ways of producing kinds or types of beliefs. It is not constitutive to the nature of any perception or perceptual belief that it is reliably veridical. A natural competence for *producing* any of these states could be relevantly unreliable. *Sustaining* such beliefs in memory depends for its warrant partly on the initial belief state that is retained. There is nothing about perceptual belief that makes it *constitutively*, for every possible believer with a perceptual system, *prima facie* knowledge. Parallel points apply to most types of empirical inductive inference.

There may be no perceptual systems in nature that are naturally unreliable in their natural habitats. There is, however, no *metaphysically necessary* bar to a perceptual system's being naturally unreliable in yielding veridical perceptions or perceptual beliefs. There are almost certainly perceptual state types in nature that are relevantly unreliable. Accuracy versus speed trade-offs are a prominent source of such cases. The competence in generating such state types is not reliably accurate in normal, natural conditions. There may be, in some individuals, whole ranges of types of perception, perhaps attributions of color shades, that are unreliable.

The key terms for evaluating (EI) are 'constitutive' and 'metaphysically necessary'. By its nature—hence constitutively—epistemic warrant marks a perceptual belief type, produced or sustained in a given way, as conducive to being knowledge. If any perceptual belief is epistemically warranted purely by virtue of being that psychological type of perceptual belief, there must be something metaphysically necessary that makes it conducive to being knowledge. I think, to the contrary, that perceptual beliefs are epistemically warranted by virtue having a metaphysically contingent property. The property is *being produced or sustained through a well-functioning belief-forming competence that reliably produces true beliefs of the same content-type by reliably connecting them to their subject matters*.

Strictly speaking, most perceptual states and perceptual beliefs are not veridical.<sup>53</sup> Human perceptual states tend, however, to be approximately accurate. The human visual system has been shown to be amazingly approximately accurate in its natural domains of application. Such "domains" include normal viewing conditions of macro-objects at relatively close range. Various optimality theorems have been proved, indicating that given such biological constraints as limitations on acuity, the human visual system is close to optimally accurate in several visual tasks, in ecologically normal circumstances.<sup>54</sup> Checks of perceptual states for accuracy use

<sup>53</sup> I discuss this issue in detail in Section 2.5.

<sup>54</sup> C. Blakemore, 'The Range and Scope of Binocular Depth Discrimination in Man', *Journal of Physiology* 211 (1970), 599–622; W. Geisler, J. Perry, B. Super, and D. Gallogly, 'Edge Co-occurrence in Natural Images Predicts Contour Grouping Performance', *Vision Research* 41 (2001), 711–24;

converging, mutually supportive tests. The conclusions about optimality are empirically grounded. They do not derive from considerations regarding the nature of perception.

There is nothing impossible about any perceptual state type's not being reliably accurate, even reliably approximately accurate. Reliability hinges on propensities of inter-relation between perceptual state types and features of the environment that they represent and that molded the natures of those perceptual state types.

Perceptual systems, and other representational systems, evolved by serving biological goods—principally, survival long enough to reproduce. Accuracy is not in itself a biological good. Nature did not select for accuracy *per se*. Accurate perceptual states evolved as part of a practically fit package. Speed and quick-trigger avoidance mechanisms can and sometimes do outweigh accuracy to yield practically valuable perception-forming routes that are not reliably accurate. Or energy saving in producing just-good-enough perception might outweigh accuracy, to yield just-fit-enough perceptual systems. It is metaphysically possible that perceptual *systems* be reliably useful without being reliably accurate. As a matter of empirically established fact, reliable approximate accuracy made enough of a contribution to fitness to be a feature of most perceptual systems, and probably most perceptual states. But no one who reflects on the evolutionary process can reasonably believe that there is a metaphysical necessity that guarantees even reliable approximate accuracy in perception. Reliable approximate accuracy is not constitutive to or metaphysically necessary for perception or perceptual belief.<sup>55</sup>

To be a perception-forming mechanism, it is enough to be embedded in a perceptual constancy.<sup>56</sup> Perceptual systems are constitutively associated, somewhere in their existence or evolution, with formation of *some* approximately accurate perceptions. The representational contents of perceptual states cannot be explained unless some of the states that are produced in the system are related, somewhere in the evolution or content-determining maintenance of the perceptual system, to perceptual states that were approximately accurate regarding environmental entities.<sup>57</sup> This point does not imply the reliability of *any* perceptual states or systems. What we know about evolution of perceptual states implies that it is possible, and sometimes actual,

D. D'Antona, J. Perry, and W. Geisler, 'Humans Make Efficient Use of Natural Image Statistics', *Journal of Vision* 13 (2013), 1–13; J. Burge and W. Geisler, 'Optimal Disparity Estimation in Natural Stereo Images', *Journal of Vision* 14 (2014), 1–18.

<sup>55</sup> Factive psychological states might be thought to be trivially reliable. One might infer that seeing-based beliefs, factively construed, are warranted by nature. Cheaply bought, cheaply made. The norm of warrant is applicable to fallible types of perceptions and perceptual beliefs. Building veridicality into the taxonomizing of psychological types bypasses one of the basic features of warrant. Warrant allows for false beliefs whose falsity hinges on incomplete information that does not result from a cognitive or perceptual defect. Warrant concerns epistemic good use of available information. Epistemic warrant fulfills a norm for our psychological capacities. Those capacities are scientifically type-individuated as fallible perceptual states and fallible perceptual beliefs. For more discussion of this issue, see my 'Disjunctivism and Perceptual Psychology' *Philosophical Topics* 33 (2005), 1–78; and 'Disjunctivism Again', *Philosophical Explorations* 14 (2011), 43–80. For further argument against taking factive states as basic for epistemology, see Ralph Wedgewood, 'Internalism Explained', *op. cit.*, 361.

<sup>56</sup> See *Origins of Objectivity*, *op. cit.*, chapter 9; and 'Perception: Where Mind Begins', *op. cit.*

<sup>57</sup> See *Origins of Objectivity*, *op. cit.*, chapter 3, esp. 67ff.

that perceptual states are not reliably accurate, even in conditions in which they evolved or were maintained. The same evolutionary considerations indicate that it is metaphysically possible that there be perceptual systems that are systematically unreliable in yielding veridicality, even in the conditions that they evolve to deal with.

So neither perceptual states nor perceptual systems are relevantly reliable *constitutively*, or by any other type of metaphysical necessity. It does not follow by necessity from what it is to be any perceptual state type that formation of instances of the type is reliably accurate, or a contributor to epistemic warrant.<sup>58</sup> An analogous point applies to perceptual systems. Since perceptual beliefs depend on perceptual states for their natures and for being warranted, and since there is no metaphysically necessary guarantee that perceptual states are reliably accurate, there is no metaphysically necessary guarantee that perceptual beliefs are reliably true.

Of course, in many counterfactual situations in which actually reliable perceptual belief types are unreliable, reflective individuals with enough experience, might develop caution that the actual individuals do not need. But (3) only depends on a metaphysical possibility. One can imagine an individual that does not get the relevant experience. Either the individual dies soon after forming the first relevant perceptual belief, or the individual is lucky in being exposed only to cases in which the perceptual beliefs are in fact true, even though their mode of production lacks a propensity to yield true belief of the relevant kind in the individual's normal environment.

Step (1) assumes that the relevant perceptual belief is warranted through entitlement. So its being prima facie warranted does not depend on background beliefs that buttress the perceptual beliefs. So the presence of either buttressing or cautionary background beliefs is irrelevant to the initial prima facie entitlement to the perceptual beliefs.

By (3), one can hold a given psychological system of states fixed, while the reliability of the fixed way of producing and sustaining a perceptual belief could vary. By (2), perceptual beliefs that are epistemically warranted in the most basic, canonical way must—as a matter of what it is to be epistemically warranted—be formed through a competence that reliably produces true beliefs. So perceptual beliefs are not epistemically warranted by metaphysical necessity.<sup>59</sup> Specific perceptual and perceptual-belief kinds could remain fixed while the epistemic warrant, the entitlement to the perceptual beliefs, could vary. So, contrary to (EI), epistemic

<sup>58</sup> A *perceptual state type*, or a *perceptual belief type*, is obtained from an instance of the type by abstracting from the occurrence of the state type in any given person at any given time, and correlatively by extracting *de re* referential applications that occur in the representational content.

<sup>59</sup> D. Smythies in 'What is the Role of Consciousness in Demonstrative Thought?', *op. cit.* and I. Dickie, in 'Visual Attention Fixes Demonstrative Reference by Eliminating Referential Luck', *op. cit.* claim that whether demonstrative reference in perceptual thought is possible depends on whether such thought has a *justification*. Nothing in the science of vision or the psychology of perception-based belief suggests any such condition on successful context-dependent perceptual reference (that is, on perception of particulars). Perceptual belief can refer to what the underlying perception can refer to. Epistemology cannot reasonably issue such armchair constraints on psychology. Neither perceptual, context-dependent reference (perceiving a particular in a given situation) nor demonstrative reference in perception-based belief depends by metaphysical necessity on any sort of warrant. To have a representational psychological competence, there must be a capacity for representational success. But the nature of the competence does not metaphysically guarantee that successes meet standards of justification or warrant.

warrant for perceptual beliefs does not supervene on, or derive purely from, the natures of the psychological states or the natures of the psychological operations that figure in the entitlement.

Perceptual beliefs' being epistemically warranted depends partly on metaphysically contingent relations between the psychology and the environment. It depends on the deeply non-accidental, but metaphysically contingent fact that reliability is biologically valuable.<sup>60</sup>

Since empirical warrant must be grounded in perceptual belief, all empirical warrant constitutively depends on factors beyond the individual's psychology. Contrary to Epistemic Internalism (EI), psychological states' being epistemically warranted does not supervene on the psychology of the individual in those states. Reliability is not coded into the natures of the states. A perceptual belief's being on a good route to truth is not insured by metaphysical necessity.

Arriving at this conclusion has depended on two types of reflection. One is, I think, *apriori*. It is reflection on the nature of epistemic warrant. There are many epistemic goods that do not require reliability. Such goods can be, and have been, confused with epistemic warrant. But the role of epistemic warrant (or "justification") in epistemic practice and the traditional role of its ancestors—Descartes's good rules for the direction of the mind and Kant's proofs of objective validity—requires that epistemic warrant mark those competencies that produce or sustain belief as conducing to knowledge. Conducing to knowledge requires that the competencies be reliable in producing or sustaining true belief. In the philosophical mainstream that ran from Descartes to mid-twentieth century, the requirement of reliability on epistemic warrant is unquestioned. Many have doubted the requirement since then. I believe that these doubts derive from a deficient understanding of the notion of epistemic warrant, and its ancestors. I have tried to show that conceptions of epistemic warrant that have been meant to motivate internalism (EI), or to doubt the reliability requirement, have straightforward defects.

The other type of reflection centers on perception. This reflection differs from anything provided in the tradition. Relevant reflection features the place of perception in evolution. The early moderns could not have reflected on evolution, of course. Philosophical work in the twentieth century can invoke no such excuse. Sense-data theory was a poorly grounded dead-end in philosophy. Post-sense-data theory of knowledge largely failed to use perceptual psychology to isolate perception as a central psychological kind. And it failed to use evolutionary biology to recognize that the reliable (approximate) veridicality of perceptual states and perceptual systems is metaphysically contingent. Since evolution does not select for veridicality, veridicality is hostage to biological usefulness. Usefulness of *reliable* veridicality in

<sup>60</sup> Natural empirical inductive/abductive inferential patterns are probably not *by nature*, or constitutively, reliable or warranted. Most *do* reliably match environmental patterns. But the reason that shows that perceptual beliefs are not constitutively warranted seems to show that empirical, non-demonstrative inferences are not either. Most of our inductive/abductive patterns probably do not constitute a logic that, purely by the form and content of the representational transitions, necessarily supports truth. This issue is, however, complex. I doubt that we know enough about empirical inductive/abductive inferences to warrant confidence on the matter. For more discussion, see 'Epistemic Warrant: Humans and Computers', *op. cit.*.

perception is not metaphysically guaranteed by any force, natural or otherwise. A consequence of these two forms of reflection is that epistemic internalism is unacceptable.

I turn to this section's final task. Recent epistemic internalism has leaned on two stock arguments. One tries to show that epistemic warrant without reason is impossible, and directly infers internalism. The other purports to show that being produced by a reliably veridical competence is not required for being epistemically warranted. I criticize these arguments.

The first argument, by Laurence Bonjour, denies warrant to a reliable clairvoyant. So reliability is not sufficient for warrant. Reliable clairvoyance is supposed to be analogous to reliable perceptual belief for which the perceiver lacks a reason-based justification. The argument concludes that reliable perceptual belief is not warranted without support by reasons.<sup>61</sup>

The argument's main point is to show that mere reliability is not sufficient for epistemic warrant. To support this conclusion, Bonjour claims that a clairvoyant would be in the rationally incoherent position of not being able to make sense of his/her own competence.

This latter claim is buttressed by suggesting that because the reliability relation is 'outside the ken' of the clairvoyant, it cannot help make the clairvoyant's beliefs warranted. It is also claimed that from the clairvoyant's subjective perspective, it is an accident that the belief is true. It is proposed that the warrantedness of the beliefs should be judged 'from the clairvoyant's perspective', rather than from a perspective that is unavailable to him/her.

As the argument claims, reliability is *not* sufficient for epistemic warrant. Not just any reliable route to truth yields epistemic warrant. The reliability must derive from a well-functioning competence that is reliably veridical, and connected in a natural, non-accidental, systematic way to the subject matters that make the warranted states or transitions representationally successful. And there must be no defeaters.

Most of us think that clairvoyance is not a real, or even nomologically possible, psychological competence. No known or plausibly conjectured mechanism connects a putative clairvoyant to relevant subject matters. Intuitions behind the argument are biased by these attitudes. If clairvoyance is to parallel perceptual belief, one must suppose that clairvoyance is a real, reliable competence—backed by some causal mechanism—and that the clairvoyant lacks reason to doubt possession of the competence. The clairvoyant would thereby be like many ordinary human perceivers, including all young children. Perhaps this supposition is impossible. Then the appeal to clairvoyance has no standing in the discussion. Let us suppose clairvoyance of this sort to be possible, for the sake of argument.

<sup>61</sup> For the argument from clairvoyance, see L. Bonjour, 'Externalist Theories of Empirical Knowledge,' *op. cit.*, 59–65. In using the argument, philosophers often move as does Bonjour, without comment, from its conclusion to epistemic internalism. However, empirical reasons are epistemic warrants. To be good reasons, they must be reliable in supporting what they are reasons for. The premises for empirical inductive arguments, and almost surely the inductive transitions from them to empirical conclusions, are not reliable by metaphysical necessity. As it stands, the argument does literally nothing to support (EI). The argument is more relevant to entitlement-to opposing step (1) of the argument set out above.

Well-functioning empirical competencies' tracking of truth is metaphysically contingent, but not accidental. To yield warrant, the competencies' reliable veridicality must be grounded in a patterned connection between (a) environmental conditions that render the beliefs true and (b) the competence's operations. So a psychological competence cannot constitutively involve help from benevolent angels, or bear a *mere* de facto correlation with its successes.

Suppose that we stipulate that clairvoyance is a psychological competence that meets these conditions. Suppose that there is a scientific explanation of the clairvoyant competence. Suppose that the clairvoyant individual lacks grounds either to believe or to doubt possession of it. Then contrary to the suggestion of the argument, despite lack of meta-beliefs about their situation, neither an individual with perceptual beliefs nor an individual with reliable clairvoyance would be in a rationally incoherent position. If one supposed that clairvoyance were a natural capacity whose reliability were grounded in real systematic relations to the subject matter, and that there were no antecedent presumptions against its viability, clairvoyants would be fair analogs to people who accept their perceptual beliefs without relying on meta-rationalizations. BonJour gives no cogent reason to believe that unsophisticated clairvoyants, under these conditions, would not be entitled to their clairvoyant beliefs. The competence's deliverances could be *prima facie* epistemically warranted.<sup>62</sup> Sophisticated individuals who knew their clairvoyance to be reliable would, of course, also not be in an incoherent position. Our background knowledge that clairvoyance is *actually* not a competence—and possibly could not be a real competence—supports our intuitive judgment that a reliable clairvoyant would lack warrant and knowledge. But our background knowledge also disqualifies actual clairvoyance as an analog to unreasoned perceptual belief.

BonJour begs the question in suggesting that since reliability is outside the ken of the clairvoyant, it cannot help make his or her beliefs warranted. Further, it would be no more an accident from the natural, competent clairvoyant's perspective that his or her beliefs are true than it is an accident from a child's perspective that his or her perceptual beliefs are true. The fact that neither the child nor the natural clairvoyant has a reason for the belief hardly makes the beliefs arbitrary or accidental to the believer.

BonJour's claim that the warrantedness of the beliefs should be judged from the clairvoyant's perspective, rather than from an unavailable perspective, again begs the question. Perceptual believers need not have meta-perspectives on their perspectives in order to be warranted. They need not have reasons. Norms for knowledge and warranted belief do apply to individuals' perspectives. They need not be applied or even applicable within those perspectives.

The other stock argument for epistemic internalism is more influential. It maintains, plausibly, that in sceptical scenarios, an individual could be epistemically warranted, but not reliable. An individual in a matrix, having proximal stimuli that

<sup>62</sup> A. Plantinga, *Warrant: The Current Debate*, *op. cit.* (Oxford: Oxford University Press, 1993) gives an example similar to BonJour's clairvoyance example to show that a reliable process need not in itself yield warrant. Plantinga notes the example's impotence against a view that sets as a *necessary condition* on knowledge a reliable well-functioning competence.



yield the same types of perceptual beliefs that would be produced ordinarily, would have warranted, unreliable beliefs. So, the argument claims, reliability is not necessary for warrant.<sup>63</sup> This argument is solemnly intoned in article after article, without any attempt to think from opponents' points of view.<sup>64</sup>

No serious philosopher who thinks that reliability is constitutive to epistemic warrant could require reliability in *all* situations. It is obvious, apart from sceptical scenarios, that there are situations in which one's warranted beliefs would be reliably mistaken. One could wander unawares into a natural situation in which illusions naturally and reliably occur. In such a situation, one would be reliably mistaken. One does not lose warrant simply by wandering into such situations. So one can be epistemically warranted in situations in which one's beliefs are reliably mistaken. The reliability that is relevant to epistemic warrant is reliability in certain normal situations.<sup>65</sup>

A contribution of 'Perceptual Entitlement', Section VI, was to specify what the normal situations are. A motive for doing so was to answer the claim that taking reliability to be a necessary condition on epistemic warrant is refuted by sceptical scenarios. The natures of all empirical beliefs are determined through interactions with the environment. The contents of perceptual states and empirical beliefs are fixed by direct or indirect causal relations to an environment. The natures of most empirical inductive inferences are also so determined. This point shows the way to determining conditions in which reliable veridicality is necessary for an entitlement to perceptual beliefs. The idea is that, for a perceptual belief to have an entitlement, the formation process that leads to the perceptual state, and from the perceptual state to the belief, must be reliably veridical *in conditions in which the perceptual state, the belief forming process, and the perceptual belief were constitutively made to have the contents that they have.*<sup>66</sup>

<sup>63</sup> As far as I know, the first arguments for epistemic internalism from sceptical scenarios occur in J. Pollock, 'Reliability and Justified Belief', *Canadian Journal of Philosophy* 14 (1984), 103–114; and S. Cohen, 'Justification and Truth', *Philosophical Studies* 46 (1984), 279–95.

<sup>64</sup> For representative examples, see R. Foley, 'What's Wrong with Reliabilism?', *The Monist* 68 (1985), 188–202; J. Pryor, 'Highlights of Recent Epistemology', *British Journal of the Philosophy of Science* 52 (2001), 95–124; R. Wedgwood, 'Internalism Explained', *op. cit.*; Matthias Steup, 'Internalist Reliabilism', *op. cit.*; M. Huemer, 'Phenomenal Conservatism and the Internalist Intuition', *op. cit.*; A. Moon, 'Three Forms of Internalism and the New Evil Demon Problem', *Episteme*, 9 (2012), 345–360.

<sup>65</sup> A. Goldman, *Epistemology and Cognition*, *op. cit.*, 107, proposes that the process for forming an epistemically warranted belief must be reliable in normal situations. There and in 'Strong and Weak Justification', *Philosophical Perspectives* 2 (1988), section VI, he sets out what I think are unpromising conceptions of what count as normal situations. In the latter article, he himself criticizes the conceptions. He abandons specifying relevant reliability as being tagged to normal situations, claiming that envatted individuals' beliefs are not warranted. I think that this view has no appeal. Specifying reliability by relation to the actual world, a move Goldman later made, is close to right. However, there are many situations in the natural world. An individual can be warranted but unreliable in some and not others. Individuals might stay longer in the unfavorable situations than in favorable situations. Such individuals' warrants do not automatically lapse. An individual's relation to the actual world can change in ways that affect warrant. Goldman's later specification is inadequate because, although it is on the right track, it is insufficiently specific.

<sup>66</sup> For a fuller account see 'Perceptual Entitlement', *op. cit.*, section VI. The following paragraphs in the text are derived from the argument pp. 531–537 of that article. Environmental conditions that sustain the perceptual content are understood to be included in normal conditions. Obviously, also, a perceptual

Here is reasoning for this conclusion.

- (I) Reliable veridicality of a perceptual-belief type—and of the underlying perceptual-state type and of the systems for forming both—are constitutively necessary conditions on an instance of the perceptual belief's having an epistemic entitlement. These requirements derive from the fact that epistemic warrant necessitates that a warranted perceptual belief be conducive to knowledge and hence on a good route to truth.
- (II) Every *prima facie* entitlement for a perceptual belief attaches to the belief via its nature or type, relative to a mode of formation or sustenance by a well-functioning psychological competence that reliably connects to the subject matter of the belief. The relevant nature of the belief is its representational content type. The force of an entitlement to a perceptual belief does not depend on the belief's relation to any other belief. No other beliefs need support it. No background considerations need protect it.<sup>67</sup>
- (III) Thus the reliability constitutively required by an epistemic entitlement to a perceptual belief is a tendency of instances of the perceptual belief's representational content type to be true, given the formation or sustaining process.
- (IV) The epistemic entitlement, hence reliable veridicality, of a type of perceptual belief, relative to the underlying formation processes, must be non-accidental. This non-accidentality must ground an explanation of why a perceptual belief of that type is conducive to knowledge and on a good route to truth. Types of perceptual belief and their formation processes are not constitutively reliable. Any type that is in fact reliable could have been unreliable. So the reliability that is predicable of a perceptual belief type must be non-accidental, but not metaphysically necessary.
- (V) What underlies this non-accidental, explanation-grounding relation between reliability and the type of perceptual belief—relative to its formation process—is the pattern of relations between the belief (and the process) and the environment that constitutively made the belief what it is. Those patterns are deeply regular, deeply systematic, and constitutively related to the perceptual belief's identity. Inasmuch as those patterns (contingently) yield a reliably veridical perceptual belief, they constitute a good, systematic connection between normal occurrences of the belief type and their being true. They constitute a deeply non-accidental basis for explaining wherein true instances of such belief types are conducive to knowledge, and *prima facie* certification of knowledge, in the senses explained in Section 2.1.

system could be formed with one set of contents; and through long-term association with another set of environmental conditions, it could take on other contents.

<sup>67</sup> If *having* a perceptual belief constitutively requires having some other beliefs, this point still stands. Having the other beliefs does not ground explanation of the entitlement to the belief, or contribute to the warranting force of the entitlement.

- (VI) Because of perspectival limitations on individuals who have perceptual beliefs—and the inevitable possibility of brute error—and because no perceptual belief is *constitutively*, by its nature, reliably veridical, no perceptual belief type can be reliably veridical under all conditions. Every type could be reliably *non-veridical* in various possible conditions that would be perceptually indiscernible to the believer.
- (VII) So warranted/entitled perceptual beliefs will be unreliable in various possible conditions other than those in which their natures—and the natures of their formation processes—were constitutively determined.
- (VIII) The only reliability relevant to the warrant/entitlement of perceptual beliefs is the reliability that obtains in conditions that ground explanation of their natures.<sup>68</sup> That is the only reliability that non-accidentally, but contingently, attaches to their types or natures in a way that grounds explanation of those beliefs' being conducive to knowledge and on a good route to truth.

The application of these ideas to sceptical scenarios is straightforward. Individuals who are moved unawares from normal environments to matrix-type situations remain entitled to their perceptual beliefs, even though they are reliably mistaken in the new environment. Being moved from one situation to another does not in itself affect epistemic warrant. Similarly, individuals born into matrix-type environments can be entitled to their perceptual beliefs. If their perceptual *systems* evolved in an ordinary environment, or if demon scientists copied stimuli from an ordinary environment, individuals' perceptual states and perceptual beliefs got their representational contents through causal relations to an ordinary environment. If perceptual beliefs would be reliable in the ordinary environment, the envatted individuals are entitled to their perceptual beliefs. Individuals in vats that bear no relation to the environmental attributes that determined the contents of *our* empirical beliefs cannot share our perceptions or beliefs.

In sum: although reliability is required for epistemic warrant, individuals lack knowledge in sceptical scenarios but can be warranted in locally unreliable perceptual beliefs.

Sceptical scenarios should not distract one from the consideration that shows reliability to be constitutively necessary for empirical epistemic warrant. The connection between being warranted and being conducive to knowledge—cited in Section 2.1 and in this section—shows that epistemic warrant requires reliability. Any belief whose formation is relevantly unreliable would, far from being conducive to knowledge, prevent knowledge.<sup>69</sup>

<sup>68</sup> Reliability in other conditions besides the content-determining conditions may be relevant to determining other entitlements. For example, a newly discovered inductive method may have to be reliable in conditions relevant to those in which it is discovered, even though no new contentful states are established. As far as I can see, however, only conditions under which the nature of a perceptual belief is fixed or sustained are relevant to determining the kind of reliability that figures in its having *prima facie* epistemic entitlement.

<sup>69</sup> James Pryor 'The Sceptic and the Dogmatist', *op. cit.*, 519 (see also 532), advocating "dogmatism", writes: 'The dogmatist about perceptual justification says that when it perceptually seems to you as if *p* is the case, you have a kind of justification for believing *p* that does not presuppose or rest on your justification for anything else, which could be cited in an argument (even an ampliative argument) for *p*.

The decisive fact that makes Epistemic Internalism unacceptable is that it offers no appropriate constitutive relation between epistemic warrant of perceptual beliefs and knowledge. Relying on certain internal matters, like perception-formation or on phenomenology, is a good thing, because they are in fact reliable. Nothing *constitutive*, or otherwise metaphysically necessary, about them makes them conducive to knowledge. What makes epistemic warrant for perceptual beliefs conducive to knowledge and a good guide to truth is a deeply non-accidental but metaphysically contingent pattern of causal relations between operation of sensory systems and the environment. These relations gave perceptual states the content that they have. They established, in the formation or maintenance of most actual perceptual beliefs, a reliably veridical relation between competence in forming those beliefs and the subject matter of the beliefs.

Internalism made sense in idealist and theological, guaranteed-harmony frameworks. For idealism, following good cognitive practice just *is* getting things right: the world just *is* the mind, operating according to rational or phenomenal rules. Idealism has shrunk to an archaic museum piece—a sign that philosophy *can* progress. Similarly, arguments for believing in a force that metaphysically guarantees reliable veridicality for perceptual states have been recognized as unacceptable. Sense-data accounts that took empirical knowledge to have an infallible base and an intrinsically reliable set of inference rules are also a dead issue. Empirical attitude contents are molded by interaction with a world that is largely independent of us. The molding is not geared to veridicality *per se*. It is geared to practicality—to fitness. No argument from the nature of perception guarantees its reliability, or hence epistemic warrant of perceptual beliefs.

Most perceptual and perceptual-belief-forming systems are reliable in their central domains of operation. Being approximately veridical at basic levels of interaction with the world does tend to serve fitness. A central feature of epistemic warrant is that it is constitutively, by nature, conducive to knowledge and a good route to truth—and, in basic cases, a *prima facie* certification of knowledge. Reliability is part of epistemic warrant's constitutive connection to knowledge and truth. Reliability is not guaranteed by anything internal. Internalism about epistemic warrant is, realistically, a lost cause.

[Note 69 Continued:]

To have this justification for believing *p*, you need only have an experience that represents *p* as being the case. Pryor adds that reliance on perceptual beliefs can be justified, even if they are not produced by a reliable belief-forming system.

Entitlement is not a dogmatist type of warrant. 'Perceptual Entitlement' *op. cit.* develops a view that is similar to the dogmatist view in holding that an individual can be warranted in perceptual beliefs without presupposing, or resting the warrant on, anything else that could be cited as an argument, or protector, for the beliefs. However, it is not the case that, to be entitled to a perceptual belief, one need *only* have an experience that represents something as being the case. The experience (or perceptual state) must derive from well-functioning, reliable perceptual and perceptual-belief-forming competencies, whose reliability is grounded in the environmental conditions in which the nature and content of the competence were determined. Pryor allows, as sufficient for warrant, that it is 'apt' from the inside to believe something. (See the text following footnote 41.) Such "aptness" *is* some a kind of cognitive good. It is not epistemic warrant. A belief formed by an "apt", unreliable system *cannot* be relevantly conducive to knowledge. By condition (c), Section 2.1, one is prevented from having knowledge if one's belief is on an "apt", unreliable route. Since epistemic warrant constitutively is conducive to knowledge and hence truth, no unreliable "apt" route could count as epistemically warranted.

### 2.3 Entitlement and Moore's Anti-sceptical Argument

A guiding idea for understanding epistemic entitlement—epistemic warrant without reason—is that perceptual beliefs are epistemically warranted, on a good route to truth, conduce to knowledge, and commonly knowledgeable. Perceptual beliefs derive from psychological transitions from perceptions. The transitions are not reason-giving propositional inferences. Perceptions are not reasons, partly because they are not propositional. Perceptual beliefs are not reasons for themselves. So perceptual beliefs are not supported by reasons. They need not be, if believers are to be warranted in holding them. Entitlement to them does not depend on support from reasons. It does not depend on background beliefs. Perceptual beliefs need only be products of a belief-forming competence that reliably produces true beliefs, through connection to their subject matters, in their normal content-determining or content-sustaining environment. The process that forms those beliefs must function well when an instance of the belief type is formed. That is the main good route to truth for perceptual belief.

This route suffices for prima facie epistemic warrant. It suffices to make true belief knowledge, in the absence of counter-warrants, internal incoherence, and Gettier conditions. To be epistemic warrant, entitlement need not borrow force from other considerations.

Warrants for perceptual beliefs are not infallible or indefeasible. Perceptual beliefs can be shown wrong by warranted theory or by other perceptual beliefs. Perceptual beliefs are basic in that warrant for them is a starting point for empirical knowledge. Normally, they have weight. Perceptual beliefs are the source of warrant for all empirical beliefs, including empirical theory. This position accords perception its natural, central place in empirical knowledge.<sup>70</sup>

The view that perceptual belief is defeasibly default-warranted, and an unaided source of knowledge, became a major position in mid-twentieth century from three sources. One was the natural idea that perceptual belief is the main basis for empirical warrant and knowledge. A second was recognition that it is a pipe-dream to think that an empirical foundational belief is infallible, or impervious to pressure from other empirical considerations. The structure of knowledge should acknowledge ordinary, fallible perceptual belief about the environment as a starting point. A third was loss of interest in tailoring epistemology to answering scepticism. Whereas Russell postulated an infallible empirical foundation to resist scepticism,

<sup>70</sup> This broad structural position is, of course, not new. Thomas Reid advances something like it: *Essays on the Intellectual Powers* (Edinburgh: Edinburgh University Press, 2002), II 5, 96; II, 22, 251–252. Mid-twentieth century versions are R. Chisholm, *Theory of Knowledge* (Englewood Cliffs, N.J.: Prentice Hall, 1st edition, 1966); A. Danto, *Analytical Philosophy of Knowledge* (Cambridge: Cambridge University Press, 1968); A. Quinton, *The Nature of Things* (London: Routledge & Kegan Paul, 1973), chapters 6–7; D. Armstrong, *Belief, Truth, and Knowledge*, *op. cit.*; J. Pollock, *Knowledge and Justification* (Princeton: Princeton University Press, 1974); W. Alston, 'Two Types of Foundationalism', *The Journal of Philosophy* 73 (1976), 165–185; 'What is Wrong with Immediate Justification?', *op. cit.*; both reprinted in Alston's *Epistemic Justification*, *op. cit.*; A. Goldman, 'What is Justified Belief?', in G. Pappas ed., *Justification and Knowledge* (Dordrecht: Reidel, 1979). I subscribe to this sort of view in 'Individualism and Self-Knowledge', *The Journal of Philosophy* 85 (1988), 649–663—see 655; reprinted in *Cognition Through Understanding*, *op. cit.*, 60.

by mid-twentieth century this strategy had come, to many, to seem not only hopeless but overwrought.

I accept these ideas. They provide some motivation for the position. They do not rule out the view that perceptual beliefs need independent apriori guarantees in order to be warranted. One might think that a perceptual belief cannot be warranted unless one has a general apriori warrant to rely on perception.<sup>71</sup> Or one might think that to protect against scepticism, one must hold that to be warranted one must have apriori warrant to believe that the scenarios are not in play.<sup>72</sup>

I think neither of these views correct. They are incompatible with our having entitlement to perceptual belief.<sup>73</sup> Entitlement derives from a reliable, well-functioning, perceptual-belief-forming system, grounded in content-determination that derives from causal interaction with the environment. No conditions on having specific background warrants need be fulfilled. Even if some individuals have such general background warrants, the warrants are not needed for one to be *prima facie* epistemically warranted in having perceptual beliefs.

Once a believer can formulate sceptical hypotheses, the believer is warranted in rejecting them out of hand, unless they are shown to be *reasonable*, not just conceivable, threats to entitlements. Sceptical hypotheses are known to be incompatible with perceptual beliefs. The individual is entitled to the beliefs. Lacking warrant of their own, sceptical hypotheses have no force to counteract an individual's entitlement. Lacking warrant of their own, they also cannot show an individual to be warranted in accepting perceptual beliefs, but unwarranted in rejecting sceptical hypotheses that the individual knows to be incompatible with those beliefs.

At least two considerations back these points.

The first is that very young children are epistemically warranted in their perceptual beliefs. They have knowledge through them. A young child knows that that body is round. Young children cannot think reasons for such beliefs. They cannot think even such relatively trivial reasons as I saw the F, because they cannot yet think reasons that make reference to their perception—much less to their belief-forming systems' being

<sup>71</sup> Examples of proponents are C. Wright, '(Anti)-Sceptics Simple and Subtle: Moore and McDowell', *Philosophy and Phenomenological Research* 65 (2002), 330–348; S. Cohen, 'Basic Knowledge and the Problem of Easy Knowledge', *Philosophy and Phenomenological Research* 65 (2002), 309–329; and 'Why Basic Knowledge is Easy Knowledge', *Philosophy and Phenomenological Research* 70 (2005), 417–430.

<sup>72</sup> Proponents of the view that we have a general apriori warrant to reject sceptical scenarios are J. Vogel, 'Cartesian Scepticism and Inference to the Best Explanation', *Journal of Philosophy* 87 (1990), 658–666; L. Bonjour, *In Defense of Pure Reason* (London: Cambridge University Press, 1998); S. Cohen, 'Contextualism and Scepticism', *Philosophical Issues* 10 (2000), 94–107; Christopher Peacocke, *The Realm of Reason* (Oxford: Oxford University Press, 2004); C. Wright, 'The Perils of Dogmatism', in S. Nuccetelli and G. Seay eds., *Themes from G. E. Moore: New Essays in Epistemology and Ethics* (Oxford: Oxford University Press, 2007). I find unpersuasive each of these claims about the content of an apriori reason for rejecting scepticism. We may have apriori reason to reject it, but such reasons are not easily found. Except for Peacocke, all these authors insist that we *need* an antecedent apriori reason, if we are to be warranted in perceptual beliefs and/or in rejecting scepticism. For an individual to *be* warranted (as opposed to just *having* warrant), the apriori warrant would, I think, have to be operative. Here I disagree with Wright. To be operative, a background warrant would have to be the content of an actual attitude. Such a requirement would leave much or all of humankind without epistemically warranted perceptual beliefs.

<sup>73</sup> They are also incompatible with how most of the authors listed in note 70 understand defeasible default warrant.

reliable. Requiring even trivial reasons as a condition on perceptual knowledge would betray a lack of epistemic perspective, or a conflation of perceptual knowledge with some higher type of knowledge.<sup>74</sup>

A second consideration is reflection on undefeated, mature human epistemic practice. Such practice does not require that individuals lean on, or even have, *reasons* for their perceptual beliefs as a condition for counting their perceptual beliefs epistemically warranted or knowledgeable. They need not lean on, or have, background beliefs that protect their perceptual beliefs. Even such simple reasons as I saw the F are glosses on an antecedent warrant. Such reasons are supplementary. They are not necessary for having knowledge. The perceptual beliefs are warranted and normally counted as knowledgeable, on their own.

When human adults understand sceptical scenarios, they normally reject them as not worth taking seriously. Scepticism dramatizes the in-principle possibility of widespread error. One need not defend against *other* merely possible errors to be warranted in relying on perceptual beliefs. No reason has been given for thinking that one must defend against the wild possibility that one is in a sceptical scenario.

Sophisticated scientists, as well as ordinary adults, commonly dismiss such scenarios. It is perhaps correct to criticize this reaction as uncurious and unphilosophical. It would be philosophical hubris to count it as not epistemically warranted. It would be hyper-intellectualized to count it warranted only because the individuals implicitly know an anti-sceptical argument, or implicitly rely on purportedly apriori warranted meta-beliefs like My perceptual beliefs are reliably formed. Sceptical scenarios have not been shown to be *warranted* threats. Defending against merely possible illusions has never been shown necessary for having perception-based warrant or knowledge. Ignoring illusions as threats, when the illusions are considered merely possible, is standard in science and ordinary thinking. Scepticism gives no reason to reject this practice.<sup>75</sup>

Sceptical scenarios are dismissed *because* they are incompatible with beliefs licensed by excellent sources of empirical warrant and knowledge: perception, perceptual belief and reasoning from perceptual belief. Claiming that the scenarios must be blocked if one is to have ordinary perception-based warrant and knowledge is wildly out of step with our best bases for forming beliefs. It is out of step without the slightest positive reason to be so.

It is a *prima facie* datum that sceptical scenarios are unwarranted. Showing *why* they are unwarranted is of real philosophical interest. Emphatically, I do not dismiss philosophical interest in scepticism. One cannot *show* scepticism wrong merely by relying on one's perceptual beliefs. It is of philosophical interest to show scepticism wrong and to explain its error. Still, no theory should count the unphilosophical scientist's disbelief in such scenarios unwarranted, even if the disbelief rests purely on the scenarios' incompatibility with perceptual beliefs.

<sup>74</sup> This point is the primary emphasis of 'Perceptual Entitlement', *op. cit.*, sections I-II.

<sup>75</sup> Many sceptical challenges implicitly depend on a sensitivity principle, to the effect that (necessarily) if A is warranted in believing that *p* (or knows that *p*), then A would not believe *p*, if *p* were false. I think that the common sense points adduced above entail rejection of any such principle. For good extensive critical discussion of the principle, see T. Williamson, *Knowledge and Its Limits*, *op. cit.*, chapter 8.

More generally, no one has shown that, for one's perceptual beliefs to be warranted or knowledgeable, one must, in the absence of a *prima facie warranted* threat, have a reason that defends against the possibility of an illusion—sceptical or ordinary. Entitlements to perceptual beliefs require no help from general beliefs. Other things equal, one can reason from warranted perceptual beliefs to rejection of propositions known to be incompatible with them.

It is a mistake to think that this response begs the question against the sceptic. The mistake lies in thinking that one *needs to* argue against the sceptic. A philosopher may want to give such an argument. Then it would not be enough just to appeal to one's entitlements. But I am not arguing against scepticism. I am not explaining the irrationality of scepticism or trying to defeat it. Those are projects for other occasions. I claim that scepticism need not be argued against, for one to be warranted in rejecting it. One can stand by one's warranted perceptual beliefs and warranted inferences, unless the sceptic gives a reason that undermines the warrant. Presentation of the mere possibility of undetectable sceptical scenarios is not such a reason.

Philosophical approaches that flout these points suffer, I think, from serious hyper-intellectualization. Tailoring an account of epistemic warrant to fit conditions of philosophical understanding is the wrong way around. Philosophy is concerned with understanding. We want to understand warrant and knowledge—place them in a wider scheme of things. In the breadth of its ambition, this ideal goes beyond what even science aims at. It is a good ideal. It is not, however, a constraint on knowledge or warrant. Ordinary, non-philosophical individuals—both scientists and the unsophisticated—are entitled to their perceptual beliefs. They are entitled to use them to dismiss sceptical scenarios. They are entitled to ignore, reject, or dismiss putative threats from in-principle possible illusions, unless they have *reason* to do otherwise. Their having these entitlements does not require having the slightest understanding of *why* they have them. Having the entitlements does not require that one understand *why* sceptical scenarios, or indeed the possibilities of ordinary illusions, are not reasonable threats.<sup>76</sup>

The tendency toward hyper-intellectualization derives partly from the ritual of initiating novices into philosophy by exposing them to scepticism. Such rituals are valuable in encouraging reflection. They should not be allowed to distort accounts of empirical warrant, by granting the mere possibilities raised by scepticism the status of threats to perceptual beliefs.

<sup>76</sup> Barry Stroud, 'Understanding Human Knowledge in General' in M. Clay and K. Lehrer eds., *Knowledge and Skepticism* (Boulder: Westview, 1989), challenges certain traditional attempts to understand knowledge in the light of scepticism. Specifically, he thinks that trying to use only materials from outside a domain questioned by a sceptic will inevitably fail. He thinks that this failure is endemic to epistemology and will inevitably leave epistemologists in an unsatisfactory position. Stroud's challenge is a good one. Still, the idea that epistemology must follow the route that he challenges, and the idea that not following it must leave one in an unsatisfactory intellectual position, seem to me unpersuasive. Philosophy is flexible about methods that it can employ to gain understanding. For further discussion of these issues, see E. Sosa, 'Philosophical Skepticism and Externalist Epistemology', *Proceedings of the Aristotelian Society, Supplementary Volumes* 68 (1994), 263–290; B. Stroud, 'Scepticism, "Externalism", and the Goal of Epistemology', *Proceedings of the Aristotelian Society, Supplementary Volumes* 68 (1994), 291–307; E. Sosa, *Reflective Knowledge* volume II (Oxford: Clarendon Press, 2009), chapters 8 and 9. Incidentally, the doctrines that Stroud and Sosa call 'externalism' are not what I count as externalism.



Some philosophers have held that we lack warrant to believe (nor do we know) that we are not in sceptical scenarios.<sup>77</sup> I think the claim purest philosophical affectation. This preposterous view derives from letting philosophical rituals distort philosophical judgment. An idea with similar pedigree is that to be warranted in rejecting sceptical scenarios, one must have a reason to block them that is independent of one's perceptual beliefs. Granting philosophy a protected space, or believing that sceptical questions have a special status that inoculates them from standard dismissive responses, is an occupational disease in philosophy.<sup>78</sup> No one has given reason to believe that sceptical challenges are special in being immune to the demand for positive rational support—not simply presentation of a possibility—if they are to undermine warrant for perceptual beliefs. It is absolutely crucial in epistemology to distinguish what is needed to understand knowledge and epistemic warrant from what is needed to have them.

The intuitions just criticized commonly underlie the idea that Moore's best-known anti-sceptical argument is no good. Here is Moore's argument adapted to our issues:

- (a) *That* [pointing] is a spherical body [warranted perceptual belief].
- (b) *That's* being a spherical body is incompatible with my being in a matrix situation with no spherical bodies that could be perceived or otherwise picked out in thought [based on logic and apriori semantics].

So (c) I am not in a matrix situation with no spherical bodies...  
[by deductive inference from the premises].<sup>79</sup>

<sup>77</sup> R. Nozick, *Philosophical Explanations* (Cambridge, Mass.: Harvard University Press, 1981).

<sup>78</sup> Contextualism about knowledge-claims is a recent instance. See K. DeRose, 'Solving the Sceptical Problem', *The Philosophical Review* 104 (1995), 1–52; D. Lewis, 'Elusive Knowledge', *Australasian Journal of Philosophy* 74 (1996), 549–567. This approach has its ancestry in ideas about relations between ordinary and philosophical questions proposed by Kant and Carnap. Clarke and Stroud pursue a similar approach, without contextualist apparatus. They distinguish ordinary "internal" from "external" questions. They take scepticism to raise a doubt to which ordinary or scientific counter-claims are irrelevant. See T. Clarke, 'The Legacy of Scepticism', *The Journal of Philosophy* 69 (1972), 754–769; B. Stroud, *The Significance of Philosophical Scepticism* (Oxford: Clarendon Press, 1984), esp. chapters 2 and 3. The analogies that they use to show that the sceptic's doubt survives standard dismissals all involve actual defeaters of purported knowledge. For example, Clarke's airplane spotters' purported knowledge is trumped by their superiors' actual knowledge that there are unusual airplanes that the spotters cannot distinguish from the more usual planes that they do spot. Sceptical doubts are disanalogous on just this point. Clarke fails to notice the disanalogy. Sceptics have no warranted or knowledgeable counterclaims. I believe that an epistemically relevant distinction between ordinary and external (scepticism-relevant) questions or contexts has never been well supported.

<sup>79</sup> Here and in some of what follows, I omit *de-re*-functioning applications on secondary attributives—as in the case of spherical in this argument. For the original argument, see G.E. Moore, 'Proof of an External World', *Proceedings of the British Academy* 25 (1939), 273–300. I express sympathy for Moore's premise and his use of it to deny scepticism in 'Individualism and Self-Knowledge', *op. cit.*, 59–60; in 'Reply to Stroud', in M. Hahn and B. Ramberg, *Reflections and Replies: Essays on the Philosophy of Tyler Burge* (Cambridge, Mass.: MIT Press, 2003); and in 'Reply to Martin Davies', in M. Frapolli and E. Romero eds., *Meaning, Basic Self-Knowledge, and Mind: Essays on Tyler Burge* (Stanford: CSLI Publications, 2003).

Others have defended Moore's argument in some detail, notably J. Pryor, 'What is Wrong with Moore's Argument?' *op. cit.* Pryor's defense of Moore's argument is skillful and insightful. Pryor rightly takes *prima facie* warrant for perceptual belief not to depend on further beliefs, or on background propositional information available in the psychology of the believer. On this point, his view is congenial to mine, and to the views cited in note 70. On the other hand, as remarked in Section 2.2 and notes 34, 41, 43, and 64 his

Moore's argument tracks adult human, including scientific, epistemic practice. The perceptual belief in the first premise, (a), has undefeated prima facie epistemic entitlement. Premise (b) is obviously true. The premises support the conclusion.

I discuss two responses to Moore's argument. Both derive from excessive concern with finding an argument that *explains* why scepticism is wrong.<sup>80</sup> One response claims that accepting Moore's argument makes knowledge "too easy". The other claims that the argument is circular.<sup>81</sup>

The warranted belief, and knowledge, obtained from the argument is "*too easy*" only relative to an expectation that a warrant must yield understanding of why scepticism is mistaken. Moore's argument does not explain why the sceptical scenario can be dismissed. It ignores or dismisses it. Normal sophisticated epistemic practice, indeed sanity, urges against accepting sceptical scenarios as actual defeaters. The argument sides with sanity without explaining why the scenarios are not justified threats to the entitlement, or why scepticism is irrational.

Warranted dismissal of sceptical scenarios is compatible with fascination about why they should be dismissed. Such fascination can motivate good philosophical inquiry. But warranted dismissal is also compatible with not caring why sceptical scenarios should be dismissed.

Similar points apply to doubters of ordinary perceptual beliefs who produce no ground for doubt. Consider 'Is it really red?', when one is looking at a surface in a good light and no special condition triggers the question. The idea that one needs to take such doubts seriously to be warranted in rejecting them is unsupported and out of keeping with standard conceptions of warrant and rationality. Such doubts are a natural focus for philosophical explanation or for teaching the young. For warrant and knowledge, such groundless doubts can be ignored.

As I have indicated, young children have warranted perceptual beliefs and knowledge without knowing what an illusion is. It seems to me wildly mistaken to believe that a young child, who cannot even think sceptical scenarios, does not know or is not epistemically entitled to the belief that *that* is red or *that* is a body. Generalized

denial that epistemic warrant (or "justification") requires reliability and his claim that epistemic warrant for perceptual belief derives from phenomenality, or from the mere fact that one has a perceptual experience that represents what the perceptual belief represents, are, I think, serious mistakes.

<sup>80</sup> Arguments that warrant is not passed from perceptual belief to the conclusion of Moorean arguments because of failure of a relevant closure principle all seem to me to be defective. For examples of anti-closure arguments, see R. Nozick, *Philosophical Explanations*, *op. cit.*; C. Wright, '(Anti)-Sceptics Simple and Subtle: Moore and McDowell', *op. cit.*; M. Davies, 'Externalism and Armchair Knowledge', in P. Boghossian and C. Peacocke eds., *New Essays on the A Priori* (Oxford: Oxford University Press, 2000); N. Silins, 'Basic Justification and the Moorean Response to the Sceptic', in T. Gendler and J. Hawthorne eds. *Oxford Studies in Epistemology* 2 (2008), 108–142. I will not canvass the large literature on this matter. I criticize some arguments for rejecting one such closure principle in 'Reply to Martin Davies', *op. cit.*. Most such arguments assume that Moorean anti-sceptical arguments are circular or beg the question. Silins' argument does not make these mistakes. His Bayesian-based argument is, however, undermined by the discussion in Section 2.5.

<sup>81</sup> The "too easy" intuition is developed in S. Cohen, 'Basic Knowledge and the Problem of Easy Knowledge', *op. cit.* and 'Why Basic Knowledge is Easy Knowledge', *op. cit.*. The circularity claim is developed by C. Wright, '(Anti)-Sceptics Simple and Subtle: Moore and McDowell', *op. cit.*; 'Warrant for Nothing (and Foundations for Free)?', *Aristotelian Society, Supplementary Volume* 78 (2004), 167–2012; and M. Davies, 'Externalism and Armchair Knowledge', *op. cit.*. Davies subsequently changed his mind.

apriori knowable propositions that really would give reasons for rejecting sceptical scenarios are hard to come by. Entitlement to perceptual beliefs does not depend on anything so grand. It normally survives groundless doubts of all kinds.<sup>82</sup>

A philosopher's gambit is to claim that although children may have warranted perceptual beliefs, philosophically initiated adults lose warrant unless they can block sceptical scenarios, or even ordinary illusions. I find such a view rationally threadbare and out of step with good epistemic practice. The mere in-principle possibility of a sceptical scenario, or an ordinary illusion, does not begin in young children as threat to an entitlement that must be addressed if the entitlement is to be retained. No reason has been given to think that it ever becomes one. Just conceiving a possible way in which one can be wrong does not undermine a warrant in ordinary cases. No reason has been given for thinking that conceiving sceptical scenarios creates an underminer of entitlements to perceptual beliefs.

Entitlements to perceptual beliefs can yield so-called easy knowledge to the effect that one is not being fooled. If an individual has concepts of illusion and of looking red, he or she can reason from an entitled, knowledgeable, though defeasible, perceptual belief as of a red surface to a warranted conclusion that the surface is not blue but lit so as to look red.<sup>83</sup> The Moorean argument that I laid out follows such a procedure. The reasoning would not *answer* a doubter. In the absence of *warranted* doubt, the reasoning is, however, correct and justified. It can constitute knowledge. The bases for the warrant and for the knowledge lie already in the perceptual system's relation to its subject matters. I return to this matter in Section 2.5.

Claims that Moore's argument is *circular* rests on assumptions that I have already criticized. One such assumption is that one's perceptual beliefs are warranted only if one has an antecedent warrant to believe that one is not under an illusion. Another is that warrant requires antecedently knowledge that perceptual beliefs are reliable. Such views hyper-intellectualize warrant and derive, I think, from living within a bubble of philosophical dialectic with scepticism.

Moore's argument *begs the question* only if it is taken to answer doubts. In fact, the argument simply exhibits a route to knowledge. Perhaps Moore took it to answer doubts. I think that he aimed just to embarrass those who think that scepticism presents a serious threat to knowledge. Moore's argument does not beg the question against sceptical doubt. It ignores such doubt. It can legitimately do so, because the doubt is not a warranted, or even supported, threat. Sceptical doubt is epistemically interesting, but epistemically insubstantial. A basic mistake in epistemology is to tailor an account of epistemic warrant to fit a response to scepticism.

<sup>82</sup> J. Pryor in 'What is Wrong with Moore's Argument?', *op. cit.* takes merely having an irrational belief that is incompatible with a perceptual belief, to undermine warrant for the perceptual belief. I see no reason to think that the irrational belief affects entitlement to the perceptual belief in any way. The irrational belief *does* make the individual's views incoherent. The individual cannot be warranted in both the perceptual belief and the irrational belief, and perhaps cannot have knowledge through the perceptual belief. However, the irrational belief does not undermine the relevant epistemic goodness of holding the perceptual belief. Since the irrational belief presents no warranted defeater, lack of warrant lies purely in it and in the incoherence of harboring both beliefs.

<sup>83</sup> S. Cohen, 'Basic Knowledge and the Problem of Easy Knowledge', *op. cit.* claims that one cannot know that one is not having an illusion in that way.

Moore's argument uses premises that are vastly more powerful than any premises that motivate scepticism. Moore makes this point in a separate meta-argument.<sup>84</sup> The meta-argument should remind philosophers that scepticism is no threat to knowledge.

For well-functioning, reliable systems of perceptual-belief formation, entitlement comes first. Defense of the entitlement must be mounted only if there is a real threat of danger—a threat with warranting force. At the heart of the structure of empirical epistemic warrant is entitlement—warrant without reason. Ordinary and scientific empirical reasoning build on that basis.<sup>85</sup>

## 2.4. Entitlement and Bootstrapping

I turn to the first of two arguments that purport to undermine the idea that we have entitlement to perceptual beliefs—prima facie epistemic warrant without reason. This argument is known as a track-record or bootstrapping argument. The argument proceeds in two stages. First, it displays an unattractive form of reasoning. Then it claims that any view that perceptual belief is prima facie warranted in a way that does not depend on reasoned support, or on warranted antecedent beliefs, is committed to the form of reasoning.

Here is the argument. Suppose that one reflects on one's perceptual state and judges,

(1') This<sub>1</sub> body appears red<sub>that2</sub>.

Suppose that one also forms the perceptual belief,

(2') This<sub>1</sub> body is red<sub>that2</sub>.

<sup>84</sup> G. E. Moore, *Some Main Problems of Philosophy* (London: George Allen & Unwin Ltd, 1953), 119ff, 141ff.

<sup>85</sup> Ernest Sosa's notion of animal knowledge bears a potentially misleading, superficial similarity to knowledge by entitlement. See his 'How to Resolve the Pyrrhonian Problematic: A Lesson from Descartes', *Philosophical Studies* 85 (1997), 229–249; *A Virtue Epistemology*, *op. cit.*, 104–105 and *Reflective Knowledge*, *op. cit.*, chapters 7–8. Roughly, for Sosa, animal knowledge is knowledge that does not require meta-reflection on belief, knowledge, reliability of resources, or the like. Animal knowledge includes first-order knowledge *based on reasoning*, hence justified knowledge, as well as unreasoned knowledgeable perceptual belief. It thus includes knowledge by entitlement, but is not equivalent with it. Sosa places meta-requirements on first-order knowledge (what is defined as animal knowledge) *in humans*, if it is to count as knowledge. These requirements entail that no humans have animal knowledge. He writes 'Human knowledge always requires a degree of meta-aptness of some minimal level. (Brute animal cognition requires a substantially lower level.)', *Knowing Full Well* (Princeton: Princeton University Press, 2011), 92–93. Meta-aptness 'requires... that the believer aptly believe, at least implicitly, in the *aptness* of his first-order belief' (94). Aptness, for Sosa, is accuracy on account of a reliable well-functioning competence. Contrary to these claims, human children have entitlement-based perceptual knowledge before they have relevant meta-competencies. They have such knowledge before they have *concepts* of belief, reliability, competence, well-functioning, or accuracy on account of a competence. The term 'implicit belief' in Sosa's requirement of meta-aptness cannot cover cases in which an individual lacks the concepts to have the belief. The requirement is not only mistaken about human children. It hyper-intellectualizes *human adult* knowledge. Even unsophisticated human adults might not understand the relevant concepts well enough to be disposed toward believing in meta-aptness. Even if they do, their entitlement-based knowledge through perceptual belief is, constitutively, independent of meta-cognitive dispositions. For criticism of similar statements of Sosa's, see J. Greco, 'Virtue, Luck and the Pyrrhonian Problematic', *Philosophical Studies*, 130 (2006), 9–34.

So (3') This<sub>1</sub> body appears red<sub>that2</sub> and is red<sub>that2</sub>.

An obvious consequence of (3'), together with the assumption that one believes (2') and (3'), is

(4') My color vision provided a true perceptual belief—the belief, this<sub>1</sub> body is red<sub>that2</sub>.

One can reason similarly regarding other bodies and other colors, each time concluding that one's color vision yields a true belief. One can then reflect on the reasoning and conclude,

(5') My perception-/perceptual-belief forming system yielded true perceptual beliefs each time: no errors were committed.

One then infers inductively,

(6') My perception-perceptual-belief forming system is reliable.

Suppose that the following view, (E), were correct.

(E) We are prima facie epistemically warranted—entitled—to such beliefs as (2'). The warrant need not be supplemented. We are entitled to (2'), independently of any antecedent warrant that we may have for believing in our reliability. It is enough, for our being prima entitled to such beliefs, that the processes that formed them function well on the particular occasion, and that the belief types and their formation types are reliable, where the reliability is grounded in the competencies' relation to the subject matter.

(7') (E) is committed to (1')–(6').

(8') Since the reasoning in (1')–(6') is unacceptable, (E) is mistaken.<sup>86</sup>

There are many moving parts in (1')–(8'). Let us go over some of them.

The reasoning in (1')–(6') does not match (E)'s account of default warrant for perceptual belief. (E) says nothing about an entitlement's making reference to beliefs or appearances ((1')–(3')), or to one's psychological system or its reliability ((4')–(6'))—except to say that beliefs about these matters is no constitutive part of individuals' being entitled to their perceptual beliefs. Young children with entitlements to their perceptual beliefs are probably not even capable of meta-beliefs about

<sup>86</sup> This sort of argument traces to R. Fumerton, *Metaepistemology and Skepticism*, *op. cit.*, 173–180. Variants are given by J. Vogel, 'Reliabilism Leveled', *The Journal of Philosophy* 97(2000), 602–623; Stewart Cohen, 'Basic Knowledge and the Problem of Easy Knowledge', *op. cit.*; R. White, 'Problems for Dogmatism', *Philosophical Studies* 131 (2006), 525–557; and S. Cohen, 'Bootstrapping, Defeasible Reasoning, and A Priori Justification', *Philosophical Perspectives*, 24 (2010), 141–159. White (553n1) takes 'Perceptual Entitlement' to be undermined by this argument and by the one discussed in Section 2.5. White mainly targets the view, 'I am justified in believing that a card is red if it appears red, even if I am not justified in believing that my color-vision is reliable. All that is required is that I have no reason to suspect that it isn't reliable'. This view is roughly James Pryor's. It is not mine. For me, more is required: reliability, grounding of reliable operations in content-determiners and truth-makers, and well-functioning of the operations. I will, however, consider the bootstrapping argument in relation to my view, (E).

their perceptions, beliefs, or psychological systems. They are not committed to (1')–(6') because they cannot think it.

Some proponents of (1')–(8') concede that (1')–(3') does not represent how (E), and other default-warrant views, regard formation of perceptual beliefs. According to (E), there is a transition from a perception as of the body's being red to a perceptual belief that the body is red. The transition is not a propositional inference. Much less does it involve a belief about one's perceptual state—about how things appear—as (1') and (3') do. The more sophisticated proponents of (1')–(8') hold that we adults can know (1') by reflection on our perceptual experiences. Once we know it, we can make such correlations as (3'). These proponents state that once one has concepts like psychological system, reliability, and well-functioning, one can reason (1')–(5'). I grant these points.

Proponents of (1')–(8') want to place the blame for the odd argument (1')–(6') on the view that (2') is default prima facie warranted. I think that the default prima facie warrant for (2') is correctly described in (E). Proponents of (1')–(8') claim that if (2') is thus warranted, the rest of the argument to (6') is inductively sound and inductively warranted.<sup>87</sup>

But the procedure is not a good induction to a conclusion about the reliability of one's perceptual belief. The problem does not reside in any view about how (2') is warranted. The form of (1')–(6') is parallel to the following. Suppose that I note (1'') I believe that Istanbul is more populous than Vienna. (2'') Istanbul is more populous than Vienna. So (3'') I believe that Istanbul is more populous than Vienna, and Istanbul is (in fact) more populous than Vienna. So (4'') My belief system has supplied me with a true belief. The procedure can be repeated for very many of my beliefs, yielding (5''): My belief system yielded true beliefs each time: no errors were committed. So (6''): My belief-forming system is reliable.

(1'')–(6'') is no better as inductive reasoning than (1')–(6'). Its failure to provide good inductive reason for the conclusion does not depend on how (2'') is warranted. I suppose that the warrant is not by entitlement. One believes (2'') from taking a census, reading travel books, asking inhabitants for population numbers, and so on. Listing beliefs that one has come to while maintaining those beliefs does nothing to support one's reliability inductively. So applying (E) to (2') is not responsible for the unattractiveness of (1')–(6'). Normal inductive evidence or argument for one's general reliability must derive from checking one's perceptual beliefs against independent indicators of their truth, from various standpoints.

In effect, (1')–(3') merely avoids a variant on Moore's paradox.<sup>88</sup> In (1')–(3'), one aligns one's belief about a perception, or perceptual appearance, with a belief that

<sup>87</sup> Cohen, 'Basic Knowledge and the Problem of Easy Knowledge', *op. cit.* believes that the induction must be accepted by anyone one who thinks that one can have warranted perceptual belief independently of having warrant that one is reliable. White, 'Problems for Dogmatism', *op. cit.*, 546, endorses the inductive inference, given that (2') is taken to be immediately warranted. White asks rhetorically, 'what else could account for my long run of successful color judgments if not the reliability of my color-vision?'. Cohen, 'Bootstrapping, Defeasible Reasoning, and A Priori Justification', *op. cit.*, 143, calls the move from (5') to (6') 'straightforward induction' and endorses White's endorsement, 144.

<sup>88</sup> G. E. Moore, 'Moore's Paradox', in T. Baldwin ed. *G. E. Moore: Selected Writings* (New York: Routledge, 1993).

directly derives from the perception. One is warranted in the matching as long as one is entitled to (2')—and as long as no counter-warrants upset one's entitlement to (2'). Both perception and belief are committal: They have a representational function whose fulfillment hinges on whether their representational contents are veridical. (See note 5.) The matching procedure consists in simultaneous acknowledgment of the commitments of the perception and the perceptual belief. Contrary to proponents of (1')–(8'), the series of matches is best *explained* not by the conclusion that one's perceptual system or belief system is reliable. It is best explained by the observation that one has aligned one's meta-judgments about one's states (whether these be perceptual states or belief states) with those committal perceptual or perceptual-belief states themselves.<sup>89</sup>

Each line in (1')–(5') is warranted, assuming that (2') is warranted.<sup>90</sup> Yet (1')–(5') does not inductively support (6'). The serial matching provides no evidence for (6'). (5') is guaranteed in advance by the Moore-like procedure. The premises are not independently generated. As an induction from the series of matches, the argument is empty.<sup>91</sup>

There is, I think, good reasoning that justifies belief in the reliability of *specific* beliefs.

Suppose that (2'), now re-labeled, is warranted. Here is the reasoning:

<sup>89</sup> There have been other diagnoses of an inductive failure in (1')–(6'). One diagnosis is that the argument rests on an inductive test that cannot be disconfirmed. This idea is discussed by White, *op. cit.*, 544ff., who credits Pryor. White endorses the claim that good inductions must risk disconfirmation, but claims that the view about warrant that he targets (for present purposes, (E)) is committed to there being an induction in (1')–(6') anyway (546–548). I think that his argument for this latter claim is unpersuasive. Incidentally, there are counterexamples to the idea that good inductions must risk disconfirmation. See B. Weatherson, 'Easy Knowledge and Other Epistemic Virtues', online <http://brian.weatherson.org/ekoev.pdf>, sourced 2014.

In any case, one would have to qualify (1')–(6') to make it strictly the case that the purported induction is not disconfirmable. Actual perceptual systems do not engender perceptual beliefs if perceptual appearances contain internal signs of fishiness. For example, a perception in which there is an appearance as of a given color may be unstable, or may yield auxiliary cues that lighting conditions are too tricky to rely upon. As a result, the probability assigned by the perceptual system to that perception type's being accurate is low. Perhaps no belief is formed on the basis of the appearance. Then an instance of the move from (1') to (2') would fail. Thus if the belief about color appearances were not paired with a belief about the corresponding color's actually occurring—and not because of background propositional defeaters, but because of internal perceptual cues—then the “test” would tend to be disconfirmed.

Since some of these internal cues can be cues to malfunction, the view that perceptual states are by their nature supportive of belief is shown here, again, to be mistaken.

Absence of internal disconfirmatory perceptual cues over a series of runs may provide, to a reflective reasoner, some slight supplementary evidence of reliability, assuming that the reasoner's perceptual beliefs *are* mostly reliable and warranted. I ignore this point in the text of this section. Absence of disconfirmatory cues cannot provide evidence by itself. Normal evidential support for reliability requires independent checks on each perceptual belief.

<sup>90</sup> This claim needs qualification. Given that we know that individuals commonly neglect defeaters, it may be reasonable to suspend belief about (4') if the list of cases is long enough. One can perhaps avoid this point by taking easy cases, and by not allowing the list to be long.

<sup>91</sup> I think that the argument does not even meet standards for enumerative induction. The premises are not appropriately independent of one another. All pairings instantiate a single type of rational commitment regarding the relation between one's meta—and first-level commitments.

(1'') This<sub>1</sub> body is red<sub>that2</sub>.

One can reflect and realize with warrant that one believes (1''). So

(2'') I believe this<sub>1</sub> body is red<sub>that2</sub>.

I think that this meta-belief is empirically warranted, since one's awareness of the occurrence of one's own sensations and perceptions is empirically warranted.

One judges correctly and with warrant,

(3'') My perceptual belief this<sub>1</sub> body is red<sub>that2</sub> is epistemically warranted.

Warrant for (3'') differs from warrants for (1'') or (2''). I offer no account of the warrant for (3''). I believe, however, that if one has no warrant to believe that there are defeaters for (1'') or for (3''), if one has the notion of warrant for perceptual belief, and if one is warranted in the first-level (here perceptual) belief (1''), one is normally warranted in believing that one's perceptual belief is warranted. At any rate, I think that, when we have a warranted perceptual belief, we *are* normally warranted in believing that it is warranted. To hold the perceptual beliefs, with warrant, and doubt that one is warranted in believing them, would normally involve some of the irrationality elicited by Moore's paradox. By the apriori warranted principle:

(4'') If a perceptual belief is warranted, the way in which it is formed or sustained is reliable and well-functioning,

I conclude, with justification,

(5'') My belief this<sub>1</sub> body is red<sub>that2</sub> is formed in a reliable, well-functioning way.

In this argument, one does not learn anything about the track record of one's perceptual-belief-forming system from (1'')–(5'').<sup>92</sup> One learns only what is implicit in one's warranted acceptance of the perceptual belief. One's reasoning is not inductive. It depends on one's actual entitlements and a general apriori principle. In each case, one simply matches one's warranted perceptual belief with a warranted belief about one's being warranted in that belief.

To get an inductive reason to believe in one's reliability, one must use other considerations. One must use such considerations as whether one is diseased, whether there are defeaters, cross-checks with other perceptual beliefs, reasoning about one's circumstances, independent indicators of the truth of one's perceptual beliefs, studies of the workings of the perceptual and perceptual-belief formation systems, and so on.

The argument (1'')–(5'') is a meta-epistemological analog of Moore's anti-sceptical argument. The argument is sound. It justifies believing its conclusion in a specific, instantiated case. But the justification rests mainly on one's entitlement to (1''). As with Moore's anti-sceptical argument, the argument is not circular. The warrant for none of its steps requires being warranted in believing that one's belief-forming mechanism is reliable. The argument can display structure in one's warrants.

<sup>92</sup> With the qualification in the last paragraph of note 89.



It does not explain, or provide evidence for, one's reliability, even in the given case. It does not answer doubts about one's reliability.

The argument form does not work for every possible individual arguer. If an individual has an unreliable or badly operating perceptual system, the individual is not warranted in accepting (1'') or (3''), and thus not in accepting (5''). One may not be able to discern one's lack of warrant from the inside. One could innocently carry out unsound reasoning.

The argument is not apriori, because not all of its premises are apriori warranted.<sup>93</sup> The argument is not independent for its warranting force from sense experience. The force of the warrant for (5'') depends on empirical entitlement to (1''). As emphasized in the previous paragraph, the argument is not shown to be sound purely through understanding it.

(6') and (7') are further errors in the bootstrapping argument. Against (7'), accepting (E) does not require accepting to (1')–(6') as an inductive argument. (1')–(6') is a bad inductive argument regardless of how (2') is warranted. Against (6'), one cannot infer inductively from instances of the Moore-like procedure that one's epistemic equipment is generally reliable.

## 2.5 Entitlement and Confirmation

In this section, I criticize a more interesting argument against the idea of epistemic warrant without reason. The argument aims to show that such warrant for perceptual beliefs is incompatible with ordinary intuitions about confirmation and with Bayesian

<sup>93</sup> S. Cohen, 'Bootstrapping, Defeasible Reasoning, and A Priori Justification', *op. cit.*, 150–155, argues that the notion of undefeated empirical default warrant for perceptual beliefs is 'incoherent'. He starts by arguing that perceptual states "provide one" with reasons for perceptual beliefs—reasons like That<sub>i</sub> body looks red, and I see that that<sub>i</sub> body is red. I think this point correct as long as two further points are kept in mind. First, such "provision" occurs only if one has meta-psychological concepts like looks and see—concepts that not all perceptual believers have. Second, the force of an entitlement to a perceptual belief does not depend on such reasons—even in adults who have such reasons. I argue both points in 'Perceptual Entitlement', *op. cit.*, sections I–II. Next, Cohen postulates defeasible suppositional inference rules of the form a looks C; so a is C. He takes reasoning by such rules to be apriori. He infers that such reasoning (from a reason like That<sub>i</sub> body looks red to a perceptual belief like That<sub>i</sub> body is red) shows that the idea of undefeated empirical default warrant for perceptual beliefs is incoherent.

The argument fails with respect to entitlement in two ways. One is that since entitlements to perceptual beliefs do not depend on reasons, no point about reasoning from propositions about how things look bears on the nature of such entitlements. The other is that it is mistaken to take the inference rule, and reasoning according to it, to be apriori. (See note 110.) By Cohen's own lights, the position that he is arguing against is incoherent only if the suppositional inference rules are apriori applicable. But where such rules provide warrant, the force of the warrant depends on reliable connection between having perceptions with certain looks and having true beliefs. This connection, expressed by the schema's 'so', is empirical. For it depends on the contributions to warrant of sensory connection to the environment, of perception, and of perceptual-belief formation. Although the perceptual entitlements do not depend on the relevant reasons, the relevant reasons depend on underlying perceptual entitlements. Entitlement is here, as often elsewhere, more basic than reason. Cohen takes his argument to tell against 'evidentialist' theories—theories that claim that empirical knowledge has its basis in evidence. I suppose that beliefs of the form a looks C can be taken to cite evidence. Then the first mistake cited in the previous paragraph would not apply. (E) is not an evidentialist account. However, most of the theories that Cohen does discuss do not construe these "evidential" reasons as basic. Cohen's argument is defective in any case, because of its assumption that the relevant inference rules are apriori.

approaches to confirmation. These issues are rather technical. However, I think that following at least the gist of the discussion should open up some important non-technical, intuitive points about how entitlement applies within actual psychologies.

Suppose that one is entitled to one's perceptual belief:

(M) that<sub>1</sub> body part is a hand<sub>that2</sub>.<sup>94</sup>

Suppose that via this warranted perceptual belief, and the obvious truth,

(A) that<sub>1</sub> body part is a hand<sub>that2</sub> only if it is not the case that I am in a sceptical scenario with a mistaken belief as of a body part and hand,

One infers:

(not-V) it is not the case that (V) I am in a sceptical scenario with a mistaken belief as of a body part and hand.

Thereby, one can know (not-V).

This is a variant of the Moorean anti-sceptical argument, (a)-(b)-(c), from Section 2.3.

There are parallel arguments for knowledgeably rejecting *ordinary* illusions. A sample conclusion is that it is not the case that I am under an illusion as of a surface's being red, whereas it is white but illuminated so as to cause illusion.

Here is the argument against entitlement—warrant without reason.

(1<sub>B</sub>) If I gain warrant for believing a proposition, my confidence in the proposition should increase.

So (2<sub>B</sub>) If (M) and (A) provide warrant for (not-V), my confidence in (not-V) should increase.

(3<sub>B</sub>) Given the sources of entitlement for (M)—the perceptual states,

(P-B) that<sub>1</sub> body part [with an attendant relevant shape, also perceptually referred to]

(P-H) that<sub>2</sub> hand

my confidence in (not-V) should decrease.<sup>95</sup>

<sup>94</sup> I have slightly altered Moore's proposition to have a content of a perceptual belief. The representational content is false (or not true) if and only if either demonstrative fails to secure a referent, or either attributive fails to be true of a referent picked out by the demonstrative that it guides, or the attributive hand fails to be true of the picked out body part.

Referential applications are representational contents whose identity is fixed by some occurrent event that functions to refer deictically. Subsequent applications of this initial application in the argument (M)-(L)-(not-V) connect anaphorically to this one. Ultimately, the applications in (M) connect to an occurrent application in the perceptual state (PH) cited below. The perception and the belief have different occurrent applications, but they are anaphorically connected. If they pick out anything, they pick out the same thing. I represent them as one application. Perceptual applications succeed in referring if and only if the perceiver actually perceives the referent. There are no such things as veridical perceptual illusions. So if the perceptual applications refer to anything, they refer to a particular that is perceived via the perceptual state that underlies the belief expressed in the first premise. For discussion, see *Origins of Objectivity*, *op. cit.*, 381–383 and Section 2.1.

<sup>95</sup> I cite two perceptions here (actually two aspects of a single perception) to allude to the complexity of transitions that lead to perceptual belief. The highest-level perceptual attribution that engenders the

Both a sceptical-scenario illusion and an ordinary illusion as of a hand could consist in the same perceptual states (P-B) and (P-H), or would at any rate be indiscernible from them. So the degree to which I suspect that that particular illusion is occurring should increase.

- (4<sub>B</sub>) The incompatibility of (3<sub>B</sub>) with (2<sub>B</sub>) undermines the idea that (P-B) and (P-H) help provide an entitlement to believe that there is no illusion—(not-V).
- (5<sub>B</sub>) One needs *independent* ground to believe (not-V) in order to be warranted in believing (M) through having perceptions (PB)-(PH).

(5<sub>B</sub>) runs contrary to the view that one can be entitled to perceptual beliefs like (M), or other beliefs formed non-inferentially from perception, without support from further beliefs or presuppositions.

Bayesian probability theory is invoked, prominently by Roger White, to yield the same result. Here I center on the Moorean argument: (M)-(A)-(not-V). The probability of (P-B)-(P-H) given (M) is greater than the probability of (PB)-(PH). The probability of (PB)-(PH) given (V) is greater than the probability of (PB)-(PH). By Bayes's Theorem, the probability of (M) given (PB)-(PH) is greater than the probability of (M); and the probability of (V) given (PB)-(PH) is greater than the probability of (V). So the probability of (not-V) given (PB)-(PH) is less than the probability of (not-V). (PB)-(PH) raise the probability of (M) and (V). Perceptual experience's raising the probability of (V) is incompatible with its giving one warrant to believe (not-V).

Here is a corollary of the foregoing argument: (not-V) follows from (M), assuming, as we are, that both (not-V) and (M) are thought by the same person in a specious present on the occurrence of the same perceptual state instance. The probability of (PB)-(PH) given (V) is approximately 1. So how high the probability of [M given (PB)-(PH)] is depends mathematically on how high the probability of (not-V) is. The probability of (M) given (PB)-(PH) is less than or equal to the probability of (not-V) given (PB)-(PH). And the probability of (not-V) given (PB)-(PH) is less than the probability of (not-V). So the probability of (M) given (PB)-(PH) is less than the probability of (not-V). So (PB)-(PH) can make one warrantably confident that (M) only if one is already warrantably confident that (not-V). One needs *independent* ground to believe (not-V) in order to be warranted in believing (M) through having perceptions (PB)-(PH). This result runs contrary to the view that one can be entitled to hold perceptual beliefs like (M)—independently of support from further beliefs, or from one's having access to further propositional considerations.<sup>96</sup>

perceptual belief (M) is (PH). (PH) is inseparable from perceptual states like (PB) and from earlier-formed perceptions as of surfaces and shapes. When a perception as of a hand occurs (see note 25), that very perception involves attribution of 3-d body part that makes up the hand, attribution of the body part's visible surfaces, shape, direction, orientation, distance, and so on. The arguments wash out this complexity, taking a single perceptual attribution to be the "justification" for the perceptual belief. This oversimplification is not crucial. It does, however, encourage ignoring probability assignments in the perceptual system that are epistemically basic to the perceptual belief.

<sup>96</sup> The foregoing reconstructs two arguments in R. White, 'Problems for Dogmatism', *op. cit.*, 531–534. Analogues occur in J. Hawthorne, *Knowledge and Lotteries* (Oxford: Oxford University Press, 2004); S. Cohen, 'Why Basic Knowledge is Easy Knowledge', *op. cit.*; and N. Silins, 'Basic Justification and the Moorean

Standard applications of Bayesian probability theory do not always coincide with intuitive judgments regarding warrant. Acquisition of beliefs against a background of lottery-like confidence assignments can yield cases in which one acquires new evidence that disconfirms a proposition even though one's subjective probability in the proposition does not go down. Similarly, one can acquire evidence that supports a proposition although one's confidence in the proposition does not go up.<sup>97</sup> John Maynard Keynes called such cases '*situations of uncertainty*'. He contrasted them with cases in which new evidence is acquired against a background of differentiated prior confidence assignments. He called the latter cases '*situations of risk*'.

Situations of uncertainty do show that the intuitions in the relevant arguments do not apply to all cases. So it is important to get right the *application* of Bayesian considerations. Some authors think that lottery-like cases or cases of extreme uncertainty are key to rejecting the Bayesian-type arguments. I think that lotteries and other situations of great ignorance before forming perceptual beliefs are *not* good models for understanding relations between entitlements to perceptual beliefs and degrees of confidence. My reasons will emerge in due course.

Here is another case in which the intuitions underlying the Bayesian-type arguments do not apply. It is plausible that we have both apriori and empirical warrant to believe arithmetical propositions. Suppose that we assign an apriori subjective probability of 1 to  $2+2=4$ . We can also gain empirical evidence for believing the proposition through counting, or through the role of arithmetic in empirical science. The subjective probability assigned to the proposition, given the empirical evidence, should be lower than 1. The overall probability assignment should remain 1. Here overall subjective probability is *not* raised by new evidence. But the new evidence does seem to add to our warrant, though at a lower probability than our initial probability.

This situation is again disanalogous to the perceptual case. The empirical warrant for believing  $2+2=4$  depends on understanding of, and belief in, the mathematical proposition. Entitlement to perceptual beliefs does not depend on capability of understanding any eternal proposition. The example does issue a relevant reminder that new evidence may not raise overall subjective probability if strong prior warrant is in place. And it again brings out that one must pay careful attention to context in applying general Bayesian updating intuitions.

The relevant context here is, of course, formation of perceptual states and perceptual beliefs. I begin by saying a little about what is known about such formation. The intuitive and Bayesian arguments depend on assigning levels of confidence to psychological states. It will be well to understand the nature of relevant states. I focus on vision and visual perceptual belief.

Response to the Skeptic', *op. cit.*. White takes the argument to undermine, among others, my account of perceptual entitlement (553n1). He focuses on Pryor's "dogmatist" view not on mine. See note 86.

<sup>97</sup> This point is well made by B. Weatherston, 'The Bayesian and the Dogmatist', *Proceedings of the Aristotelian Society* 107 (2007), 169–185; P. Kung, 'On Having No Reason: Dogmatism and Bayesian Confirmation', *Synthese* 177 (2010), 1–17; and E. Zardini, 'Confirming the Less Likely, Discovering the Unknown' in D. Dodd and E. Zardini eds., *Scepticism and Perceptual Justification* (Oxford, Oxford University Press, 2014). All cite Keynes as having made the point first.

In formation of visual perceptual states, the proximal stimulus—the light array just before it hits the retinal receptors—is the *retinal image*. The *initial sensory registration* of the retinal image is the first-occurring informational state of the perceiver on which the formation of perceptual states depends. After that initial registration, a series of transformations leads to a *perceptual state*—say, a visual perceptual state that represents as of a surface’s being at a certain distance. Further transitions from this perceptual state lead to others. For example, a perceptual state as of a surface might lead to a perceptual state as of a three-dimensional body. This latter perception might engender a further perceptual state as of a body part and hand (see note 25). In some perceivers, many perceptual states transition to *perceptual beliefs*. There is a transition from a singular perception as of a surface (functioning to pick out a particular surface) to a belief that *that* is a surface. There is a transition from a singular perception as of a hand to a belief that *that* is a hand.

The three primary types of state here—initial sensory registration, perceptual state, perceptual belief—have different formats.

The initial sensory registration carries information about the proximal stimulus—the array of light intensities that strike the retina. The initial sensory registration is not a perception and does not by its nature have veridicality or accuracy conditions. The sensory registration is caused by, correlates with, and functions to correlate with the proximal stimulus. It is the basis for perceptual processing and all subjective probability assignments. Still, subjective probabilities regarding *accuracy* or *truth* apply only to perceptual states and beliefs, not to registration of retinal information.<sup>98</sup>

Second, perceptual states have veridicality conditions for *accuracy*, as an aspect of the psychological kinds that they are. They represent non-propositionally.<sup>99</sup> My formulations of (PB) and (PH) are non-propositional. There is a transition from (PB) to (PH) in the perceptual system. Assuming, as I am, that there are states produced by a perceptual system that are as of hands (see note 25), there is a transition from the perception as of a certain shaped body part (PB) to a perception as of a hand (PH). Antecedent perceptual states led to (PB). Perceptions as of bodies and body parts are formed from perceptions as of shapes, textures, motion, surfaces. This broad order of formation is well-known in perceptual psychology. There are

<sup>98</sup> For discussion of lack of veridicality conditions in sensory registrations, see *Origins of Objectivity*, *op. cit.*, chapter 8. Despite the fact that subjective probability regarding *veridicality* does not apply to information registrations, Bayesian assignments of prior and conditional probabilities of *occurrences* of such registrations do apply.

<sup>99</sup> Work in perceptual psychology does not treat perceptual states as having propositional form. I believe that this treatment is correct. The representational function of perception is to identify particulars, using attributives entirely to serve identification. Function determines form. The representational contents of perceptual states have the form of singular representations guided by attributives. The form is analogous to that of a noun phrase governed by a context-dependent determiner applied on a particular occasion. All perceptual attributives for perceivable particulars accompany applications that function to refer to instances of relevant attributes. The form differs from that of a noun phrase in language in that it is iconic. The iconic format of visual perceptual representational contents still, however, includes contextually applied determiners and attributives. For further discussion, see *Origins of Objectivity*, *op. cit.*, 537–546 and ‘Origins of Perception’, *Disputatio* 4 (2011), 1–38; ‘Steps Toward Origins of Propositional Thought’, *Disputatio* 4 (2011), 39–67; also Elizabeth Camp, ‘Thinking with Maps’, *Philosophical Perspectives* 21 (2007), 145–182. Further discussion occurs in *Perception: First Form of Mind*, forthcoming. See note 21.

also transitions from perceptions to beliefs. Less is known about these transitions. I assume that they tend to be automatic and direct.

Third, the representational contents of perceptual beliefs are propositional. Recall from Section 2.1 that *perceptual beliefs* take over perceptual attributives and embed counterparts in a propositional structure. They include occurrent applications of demonstrative-like determiners that take over, anaphorically, referents of occurrent applications of the perceptual determiners.

Proponents and opponents of the intuitive and Bayesian arguments consistently ignore basic features of perception and perceptual belief. This practice gives the arguments, and discussions of them, an overly abstract, almost make-believe character.

I cite two facts about the formation of perceptual states and perceptual beliefs that the relevant arguments commonly ignore.

One is that there is a change in representational format from non-propositional perceptual state to propositional perceptual belief. Some forms of the arguments take (M) to be justified by a further *belief* about a perceptual appearance. Passing over the distinctive way in which perceptions figure in yielding warrant for perceptual beliefs raises questions about whether the arguments apply to the epistemology that the arguments purportedly target.

Ignoring differences in format and function between perceptual states and perceptual beliefs runs together justification and entitlement.<sup>100</sup> Entitlement to perceptual beliefs is grounded in non-propositional transitions within the perceptual system and between the perceptual system and perceptual belief. Epistemic warrant attaches only to propositional attitudes and propositional inferences. States warranted through entitlement are starting points in the structure of empirical epistemic warrant. States that are empirically warranted through justification are always warranted either through self-evidence or from inference from other warranted propositional states. Blurring the distinction between propositional attitudes and perceptual states tends to distort understanding of base-line conditions for entitlements to perceptual beliefs.

<sup>100</sup> White not only takes the content of the *warrant* for (M) to be propositional. He takes it to be a belief about perceptual appearance: the belief it appears to me that there is a hand, 'Problems for Dogmatism', *op. cit.*. At 534–535, he remarks that in taking such a meta-belief as the initial warranting state, he over-intellectualizes the actual situation. He notes that in being warranted in a perceptual belief, one need not reflect (I add, or even be able to reflect) on the relevant perceptual state. He claims, however, that the transposition does not matter. He writes, I think plausibly, '... it surely should not make any significant difference whether I choose to consider how things appear to me, and form the belief that this appears to be a hand, or I don't consider the matter. If the rational response to its appearing that this is a hand, when I also believe that it appears that this is a hand, is to decrease my confidence that it is a fake-hand, then surely this is the rational response to the same experience when I do not even consider how things appear to me.' What I think plausible here is that whatever epistemic force the judgment about the appearance has in favor of believing (M) derives mainly from the formation of the perceptual state that the judgment is about. One can let the appearance judgment, to which one is entitled, count as a reason for (M). The reason is a proxy for the epistemic force present in the entitlement to the perceptual belief. While I do not think that White's transposition is the crucial error in the argument against an entitlement-based account of warrant for perceptual belief, I do think that it obscures the difference between how things work in a sub-propositional perceptual system, and among perceptual beliefs, on one hand, and how individuals operate with background knowledge, on the other. These points are elaborated below.

I think, however, that the Bayesian arguments can be formulated in light of differences between perceptual states and perceptual belief, so as to remain *prima facie* threatening. I have tried to do so above. Subjective probabilities for veridicality apply, both intuitively and in science, to both propositional states—the beliefs—and to non-propositional states—the perceptions.<sup>101</sup>

In perceptual psychology, perceptual processing, which involves no propositional attitudes in the perceiver, is often modeled in a Bayesian framework.<sup>102</sup> Bayesian subjective probabilities are assigned to perceptual states with no assumption that the states are propositional.<sup>103</sup> Perceptual states and beliefs are *committal*: success and failure of their representational functions depends on whether or not the states are veridical. Subjective probability assignments for confidence in veridicality make sense with respect to any committal state in a psychology. Subjective probabilities that are assigned to perceptual states are not conscious and do not result from individually controlled acts. They are embedded in the operations of the perceptual system. They model degrees of confidence, reliance, or anticipation. Terms like ‘credence’, which are essentially associated with belief, are inappropriate. The assignments are *subjective* in that they bear on the perspective of the individual or the individual’s states, sometimes concern the accuracy of such states, and may fail to coincide with objective probabilities.

<sup>101</sup> L. Moretti, ‘In Defense of Dogmatism’, *Philosophical Studies* 172 (2015), 261–282, argues that White’s argument fails because of its assimilation of perception to a judgment about appearance. As note 100 indicates, I do not agree. (Incidentally, Moretti treats perceptual content as propositional and, like White, does not note the *de re* applied character of perceptual, or perceptual-belief-, content.) His position hinges on claiming (273–276) that the appearance judgment lacks the warranting force of the relevant perception, and that when one makes the judgment, one loses warrant for perceptual beliefs. That is, one cannot with warrant move from the appearance judgment to the perceptual belief. This claim is, I think, implausible and unjustified. In fact, the transition from the appearance judgment to the perceptual belief is a good one *because* (and only because) the transition from the perception to the belief is a good one. Moretti’s position is of a piece with the philosopher’s gambit criticized at the end of Section 2.3. If the position were correct, it would provide a pyrrhic victory for the view that perceptual beliefs are warranted through perception, unaided by background propositional resources. One’s non-reflective warrants would be undermined by the Bayesian considerations, as soon as one reflected on one’s perceptual states. The position also does not explain well why the Bayesian arguments cannot be reformulated to apply to warrant for perceptual belief through perception. Moretti does make some sound critical points against other defenses of “dogmatism”.

<sup>102</sup> See A. Yuille and H. Bulthoff, ‘Bayesian Decision Theory and Psychophysics’ in D. Knill and W. Richards eds. *Perception as Bayesian Inference* (Cambridge: Cambridge University Press, 1996); W. Geisler and R. Diehl, ‘A Bayesian Approach to the Evolution of Perceptual and Cognitive Systems’, *Cognitive Science* 27 (2003), 379–402; W. Geisler, J. Perry, B. Super, and D. Gallogly, ‘Edge Co-occurrence in Natural Images Predicts Contour Grouping Performance’, *op. cit.*; L. Maloney and P. Mamassian, ‘Bayesian Decision Theory as a Model of Human Perception: Testing Bayesian Transfer’, *Visual Neuroscience* 26 (2009), 147–155; J. Frisby and J. V. Stone, *Seeing* (Cambridge, Mass.: MIT Press, 2010), esp. 308–322; J. Burge and W. Geisler, ‘Optimal Disparity Estimation in Natural Stereo Images’, *op. cit.*. Similar Bayesian methods have been used in sensori-motor psychology. See R. Shadmehr and S. Mussa-Ivaldi, *Biological Learning and Control* (Cambridge, Mass.: MIT Press, 2012), chapters 5, 10–12; D. Wolpert and M. Landy, ‘Motor Control is Decision-Making’, *Current Opinion in Neurobiology*, 22 (2012), 1–8. For good overviews, see M. Rescorla, ‘Bayesian Perceptual Psychology’, *The Oxford Handbook of the Philosophy of Perception*, M. Matthen ed. (Oxford: Oxford University Press, 2015); and ‘Bayesian Sensorimotor Psychology’, *Mind and Language* 31 (2016), 3–36.

<sup>103</sup> The science tends to focus on states, as do I. Some accounts focus on transitions.

Subjective probabilities assigned in a perceptual system can be discovered by inferring from reports, if the perceiver has language. Discovery need not depend on reports. It can center on action patterns—including hesitancy and speed of action—and in interactions among perceptual systems. For example, perceptual systems make compromises attuned to the relative probabilities of accuracy among perceptual states in different modalities (say, vision and touch).<sup>104</sup> Perceptual systems factor in noise, acuity, distance, lighting conditions, and so on, in determining the degree to which given types of perceptual states are relied upon for being accurate.

Any epistemology that connects with the actual psychology of perception and perceptual-belief formation must accommodate scientific ways of describing these matters. Bayesian descriptions are not ubiquitous in perceptual psychology. However, they are prominent enough that a reasonable epistemology should accommodate them. The Bayesian arguments in current epistemology do not connect to Bayesian accounts in perceptual psychology at all.

A second feature of perception and perceptual belief is the singular, occurrent, *de re* character of their representational contents. This feature is never remarked. Relevant representational contents are treated not only as propositional, but as eternal. This feature is the most important one to bear in mind if one is to understand wherein the arguments mislead.

I express the representational contents of (P-B), (P-H), and (M) with subscripts on the demonstratives. The subscripts mark occurrent events at particular times—the actual occurrence of a perceptual state or instantiation of a perceptual belief. These occurrent events are always present in the representational contents of perceptual states and perceptual beliefs. The events are *applications*.<sup>105</sup> Applications function to refer—to apply the guiding attributives, like hand, to particular instances of the indicated attribute. The function is realized successfully when a particular is perceived and the perceptual attribution to it is accurate of it. Here, the function is realized successfully when the application perceptually refers to a particular hand.

The *de re* character of perception and perceptual belief is pervasive. Every property—and kind-attributive in perception and perceptual belief guides and is accompanied by an application that functions to be *de re*. (See note 26.) Of course, in an actual visual perception of and as of a hand, one perceives not only a hand, but the constituting body part, the instance of its shape, the instance of its color, the instance of its texture, the instances of its surface parts—all *de re*. The perceptual beliefs to which we are entitled are similarly pervaded by applications that function to be *de re*. Any property- or kind-attributive in a belief that lacks such an application is not being applied in perception. It might rather be the product of association in long-term memory, inference, supposition, interlocution, or the like. So to understand the empirical starting points for empirical theory, one must take the representational contents of perceptual beliefs not only to contain applications that function to be

<sup>104</sup> M. Ernst and M. Banks, 'Humans Integrate Visual and Haptic Information in a Statistically Optimal Fashion', *Nature* 415 (2002), 429–433.

<sup>105</sup> Strictly, the same application type can occur more than once, in different events at different times, if all the occurrences are anaphorically related to a single occurrence that is the primary event of referring. See earlier discussion of applications in Section 2.1 and notes 24, 26, 59, and 94.



*de re*. One must regard every property- or kind-attributive to be referentially applied in a *de re* manner to a purportedly perceived particular.

I emphasize that a perception as of a hand cannot occur separately from a cluster of accompanying perceptual attributions. Here perception contrasts with linguistic attribution. Linguistically, attribution of *hand* can be separated from attributions of hand-shape, hand-size, and so on. A perceptual attribution as of a hand that is part of a perceptual attribution as of a [hand-shaped] body part with a certain orientation and size may be assigned a different subjective probability than that assigned to perceptual attribution as of a hand that involves different attributives for bodily shape and size. A perception as of a hand as at a great distance or as partly occluded will receive a different subjective probability than a perception as of a hand that is close up and unoccluded. So there are many ways of perceiving something as a hand—and perceptually believing something to be a hand.

The *de re* nature of perceptual states and perceptual beliefs forces a distinction between a kind or type of perceptual state and an actual occurrence of a perceptual state. The kind or type is marked by the schematic representational content without the applications. (See note 58.) It is like ‘that hand’ without an occurrent application of ‘that’. An occurrence of a perceptual state or perceptual belief always has applications in the representational content. This distinction yields structure in assignments of subjective probabilities that is missing for eternal, non-*de-re* contents. Different probability assignments attach to the *occurrence* of an instance of a perceptual state type and the *veridicality* of occurrences of perceptual state types. Occurrence and veridicality are importantly different matters.

I turn to fuller discussion of formation of perceptual states, with an eye to its connection to confirmation. Patterns of perceptual-state generation were formed through eons of sensory interaction with the environment. For example, light patterns typical of borders of shadows differ from light patterns typical of edges of bodies or surface reflectances. Differences between darker and lighter light patterns that mark shadows’ edges tend to be fuzzy. Differences that mark a body’s edge or a change in surface reflectance tend to be sharp. Exceptions occur, but are rare in natural scenes. Perceptual-state formation capitalizes on this pattern. Perceptual systems tend to discount shadows in determining edges of bodies or edges of surface color boundaries. In effect, the environmental regularity is coded into laws governing perceptual state formation.

Formation laws closely fit statistical distributions of properties in the environment. If one value of a scalar property is more common than another, and the property is visually important, such relative frequencies are likely to be encoded in subjective probabilities in visual processing.

The development of such probabilities derives both from evolutionary selection and from feedback in individuals’ learning histories. Some subjective probabilities are innately coded, through selection, in perceptual systems. At any given time in the life of a perceiver, such innate codings will have been adjusted through feedback—adaptation, priming, learning, and so on.

As indicated earlier, even though they do not have propositional structures, perceptual states are committal and have accuracy conditions. Many mainstream

theories in visual psychology treat perceptual states as having subjective probabilities for accuracy.

The handling of accuracy and inaccuracy in visual psychology is much more nuanced than its handling in colloquial speech. In ordinary speech, we commonly take our perception of distances, at least at close ranges, to be accurate. In the science, where relatively exact measurements are made, strictly accurate perceptual attributions are relatively rare, for scalar attributes. A surface, even at a distance of 100 cm, is not likely to be perceived as *exactly* at that distance, largely because of noise in the system. Visual systems can discriminate very small, fraction-of-a-centimeter differences at close ranges. Errors, evinced in the course of actions, are corrected through constant feedback.

Colloquially, one thinks of a single subjective probability for accuracy and a single one for inaccuracy, with the probability for accuracy relatively high. Simple yes-no perceptual attributions, such as whether a surface is foreground or background, will tend to be strictly accurate. With most attributions—for example, attributions of scalar attributes like distance, length, color, shape, and size—the objective and subjective probabilities for strict accuracy will be relatively low. The science focuses on a distribution of different errors, with different subjective probabilities for each. Inaccuracy might derive from malfunction, noise, limited acuity, adaptation, priming, interference from another system, or brute error.<sup>106</sup>

So the objective and subjective probabilities of getting the distance of a surface exactly right are commonly low. Commonly, getting the distance exactly right is at the top of a bell curve, with slight errors grouped around it, and larger errors falling off down the curve with sharply decreasing probabilities. There is a range of objective and subjective probabilities for types of error, including different magnitudes of error. Higher probabilities tend to be assigned in perceptual systems for small errors—such as attributing distances of 99.95 cm to a surface that is 100 cm away. Lower probabilities tend to be assigned for large errors, such as attributing distances of 400 cm, or having an illusion as of there being a surface at all.<sup>107</sup>

There is another layer of complexity in the science's account of assignments of subjective probabilities in perceptual systems. Let us focus on a visual attribution of the distance of a surface. Given a registration of a retinal image, the system is modeled as producing a distribution of probabilities for different distances: 100 cm probability of .01; 100.05 cm/.009; 99.95 cm/.0095; . . . . The system then applies a cost function for the different degrees of errors, and an estimator function that selects a distance for the perceptual state to attribute—say 100 cm. It is natural to assume that the distance selected is paired with the highest probability. Because of anticipated costs for different types of error, the system cannot be expected always to select the distance attribution that has the highest probability of being strictly accurate. I will

<sup>106</sup> *Brute error* is error that derives from a well-working psychological system that happens to be in a situation that yields the error. Brute error can be caused by residual noise. The notion is discussed in 'Perceptual Entitlement', *op. cit.*, 507ff.. See the text above associated with note 12.

<sup>107</sup> I do not assume that perceptual attributions as of distance involve measurement units like centimeters. They are likely to be unit-free magnitudes. For a fine discussion of this notion, see C. Peacocke, 'Magnitudes: Metaphysics, Explanation, and Perception', in *Mind, Language and Action: Proceedings of the 36th International Wittgenstein Symposium* ed. D. Moyal-Sharrock, V. Munz, and A. Coliva (Berlin: de Gruyter, 2015).

count as *approximately accurate* any perceptual attribution that is either strictly accurate or a near miss.<sup>108</sup>

The science's assumption that, for scalar properties, strictly accurate perceptual attributions are not common provides no encouragement to philosophical error theories about perception. The objective probabilities of small errors that cluster around the correct attribution are in aggregate very high. The objective probabilities of serious mis-attributions are low. Optimality theorems (see note 54) show that, in cases so far tested, visual systems make the size of errors close to as small as is theoretically possible, given the systems' natural limitations—such as noise and acuity. Subjective probability distributions for errors are close to objective probability distributions. Individuals' degree of reliance on one modality (say, vision) in relation to another (say, touch or audition) matches the relative objective probabilities for accuracy of the different modalities in contexts where their representations of a situation conflict. (See note 104.)

So it is appropriate to say that although most types of visual attribution tend to be strictly inaccurate, most tend to be approximately accurate. They tend to be as close to accurate as physical limitations allow. It is objectively probable that deviations from accuracy are small. Perception yields approximately true and approximately knowledgeable perceptual beliefs in many important cases. Given evolutionary time, reliance on a given type of visual perceptual state has come for the most part to closely match the objective probability of that type's being accurate.

In subsequent discussion, I write of *approximate veridicality* and *radical non-veridicality*. Approximate veridicality includes both strict veridicality and relatively small non-veridicalities. Radical non-veridicality includes large non-veridicalities and referential illusions, such as whether there is a surface at all at an attributed distance. I assume that every occurrence of a perceptual state or perceptual belief is either approximately veridical or radically non-veridical. When philosophers discuss sceptical and ordinary illusions, they are large errors, not near misses. My simplification will not, I think, materially affect the discussion.

Subjective probability assignments for veridicality, for small non-veridicalities, for large non-veridicalities, and for referential illusions *are all—when warranted—warranted empirically, not a priori*.<sup>109</sup> Even assignments encoded innately in a perceptual or perceptual-belief-forming system are empirically warranted. What warrant those assignments have depends for its force on reliable sensory contact to relevant entities in the environment, in the system's evolution. A warrant's force is a priori if that force does not derive from sensory input. Whether a state—or

<sup>108</sup> For discussions of cost and estimator functions, see A. Yuille and H. Bulthoff, 'Bayesian Decision Theory and Psychophysics', *op. cit.*; W. Geisler, 'Ideal Observer Analysis' in L. Chalupa and J. Werner eds., *The Visual Neurosciences* (Boston: MIT press, 2003); M. Banks, J. Burge, and R. Held, 'The Statistical Relationship between Depth, Visual Cues, and Human Perception', in M. Landy ed. *Sensory Cue Integration* (Oxford: Oxford University Press, 2011).

<sup>109</sup> Although, epistemic warrant attaches to perceptual belief, not to perception, I allow myself here to write of warrant for perceptual states. Perceptual states have an analog of epistemic warrant. Warrant for perceptual beliefs depends mainly on the well-functioning and reliability of the perceptual states. What I write here about sub-epistemic "warrant" for perceptual states carries over to the genuine epistemic warrant (entitlement) for perceptual beliefs.

probability assignment—occurs in an individual's psychology before the individual has sensory input is irrelevant.<sup>110</sup>

Since the discovery of evolution, it has been clear that being determined in the mind before an *individual's* sensory contact with the world is not being determined independently of sensory contact in the individual's psychological systems' evolutionary history. The basic types of perceptual attributives and basic forms of perceptual processing are set through interaction with the environment before the birth of any human individual. The degree to which a given type of perceptual state is likely to be accurate depends heavily on patterns of sensory interaction with the environment in the evolutionary pre-history of the individual. The force of one's warrant for perceptual beliefs depends primarily on the reliability of perceptual accuracy. To take assignments of subjective probabilities to perceptual beliefs to be apriori *because* they are determined before an individual's experience occurs would be to deem many of our specific perceptual beliefs apriori warranted—an absurd result. All warranted subjective probability assignments for specific perceptual beliefs that are discussed here are empirically warranted.

Subjective probabilities are assigned in the first instance, for a perceiver and time, to *types* of perceptual state. Recall that actual perceptual states always have occurrent, *de re* elements—applications. State *types* are representational contents of a perceptual state, for a given perceptual modality, with occurrent applications removed. (See note 58.) For an individual at a time, all possible occurrences of a given perceptual state type receive the *same* subjective probability assignment, if they derive from the same antecedent state types.<sup>111</sup> Probability assignments concern the role of *types* of states in *patterns* in an individual's psychology, and between the psychology and the environment. Few perceptual-state types occur at any given time. The system attaches subjective probabilities directly not to occurrent states, but to all and any occurrences of state types. To engage in a standard idealization: at any given time, every perceptual state type that could be produced in an individual's perceptual system is assigned a subjective probability for approximate accuracy and for radical inaccuracy—strictly, for strict accuracy and for types and degrees of inaccuracy.

The subjective probability for veridicality assigned to a perceptual state depends on the success record, in evolutionary history and as recorded in feedback during

<sup>110</sup> Some have held that if a probability can be determined *before* an individual's sensory contact with the world, that makes it apriori. Such views exhibit serious misunderstanding of what apriority is. Apriority has always concerned the nature of a warrant, even in pre-Kantian apriority concepts. See my 'Frege on Apriority' in *New Essays on the A Priori*, C. Peacocke and P. Boghossian eds., (Oxford: Oxford University Press, 2000); reprinted in *Truth, Thought, Reason: Essays on Frege* (Oxford: Oxford University Press, 2005). Before Kant, some ran together psychological priority, particularly innateness, and apriority, because they thought that innateness entailed independence of the force of a warrant from experience. Leibniz knew better. Kant, the father of modern conceptions of apriority, knew better. He stressed a distinction between psychological origin and justificational priority. See *Critique of Pure Reason* A1/B1.

<sup>111</sup> For a given individual and time, there could be *different* applications of the same content. In the actual world, applications are different if they apply to different particulars. An individual in the actual world could be presented with perceptually indiscernible particulars on different occasions. Then the individual's perception—and belief-types involve different applications.

the perceiver's life, of that type of perceptual state. The same perceptual-state types in different individuals, with different sensory histories, usually have different subjective probability assignments. Different types within an individual tend to have different success records, hence different assignments. As noted, a visual system tends to assign a higher subjective probability for being approximately accurate to a perception as of a close surface than a perception as of a distant one. Another way of phrasing this point is that, built into the laws of a visual system are anticipations that attributions of close surfaces are likely to make smaller errors than attributions of distant surfaces. A huge number of probability assignments are present in a perceptual system at any given time—*antecedent* to the acquisition of an occurrent perception. A huge amount of information about subjective probability of approximate accuracy and radical inaccuracy is modeled as coded into a perceptual system at any given time. Such coding is in place before the system receives stimulation that causes an instance of a perceptual state type.

Different subjective probability assignments for accuracy and inaccuracy of perceptual state types apply at each level of processing. Perceptual states formed at later stages depend on states at earlier stages. Glitches are possible at every transition. So the subjective probability of the approximate accuracy of a perceptual state type at a later stage is often slightly lower than that of a perceptual state type at an earlier stage. A system normally assigns a very slightly higher subjective probability for the approximate accuracy of a perception as of a 3-d shape than to that of a perception as of a body with that shape. It normally assigns a higher probability to the perception as of a 3-d shaped body part than to a perception as of a hand with that body-part shape.

The points made about subjective probability assignments to perceptions carry over to those for perceptual beliefs. Like perceptual states, perceptual beliefs function to be *de re*. Perceptual beliefs are formed as direct conceptualizations of perceptions. For any given believer and time, all possible instances of a specific *type* of perceptual belief, with possibly different occurrent applications, receive *the same* distribution of subjective probability assignments for strict accuracy and for various types of errors. On our idealization, each perceptual belief type that can be generated in a perceptual-belief-forming system receives, for a given believer and time, a distribution of subjective probability assignments. The distribution aggregates, I am supposing, into a dual assignment for approximate truth and radical error. Those two assignments add to 1. Normally, the main determiner of the subjective probability for approximate truth of a perceptual belief is the subjective probability for approximate accuracy of the perceptual state that the belief conceptualizes. The perceptual belief's probability is normally very slightly lower—to allow for glitches in the transition from perception to belief.

To fix ideas, I sketch a Bayesian model for a specific case. I begin by discussing *objective probabilities*. The environment contains attributes—kinds, properties, and relations—that the visual system functions to represent and attribute. Any given perceptual representation functions to represent a scene made up of instances of many attributes. I simplify, fixing on one attribute: the *distance of a surface*, symbolized by the variable: *D*. The psychologist, following in the footsteps of evolution, records the statistical distribution of various distances of surfaces from the perceiver's

eyes at any given time. The nearest visible surfaces are several centimeters away. The farthest are at the horizon.<sup>112</sup> The probability distribution peaks at the most probable distance of a surface within this range and falls off from that peak. That is, there is a probability distribution that integrates to 1 that gives the objective probability that each discernible surface distance will occur in relation to the perceiver. Focusing here on objective probabilities,  $p(D)$  is the objective *prior probability* that distance  $D$ , for each  $D$ , is a surface's distance from the perceiver.

Environmental attributes cause proximal stimuli, which cause initial registrations of the proximal stimuli. For vision, initial registrations are of the retinal image. I symbolize the initial registration of a proximal stimulus  $I$ . There is a distribution of objective prior probabilities for the occurrence of the various possible initial registrations. As an idealization, for each possible initial registration, there is an objective prior probability that it will occur. I symbolize the objective probability of the occurrence of  $I$ :  $p(I)$ .  $p(I)$  is a distribution of probabilities. For a particular value assigned to  $I$ , I write ' $I_r$ '. ' $I_r$ ' designates any member of a class of initial registrations that do not differ in ways that affect a perceptual state: all produce the same perceptual state (not merely perceptual states that represent the same attributes). So small differences among initial registrations that are within some criterion for unimportant differences are ignored in use of the symbol ' $I_r$ '.  $p(I_r)$  is a specific probability within the probability distribution. One could regard ' $I_r$ ' as a canonical name. Then  $p(I_r)$  would be a specific probability within the probability distribution.

For scenes in which the stimulus is not impoverished and noise in the system is ordinary, some probabilities that given initial registrations will occur, for a given distance value  $D_j$ , will be higher than others. The probabilities are found by studying environmental properties, optics, and properties of sensory receptors. The science has done a lot of work measuring such probabilities.

<sup>112</sup> For detailed discussion of measurement of the statistics of various attributes in natural scenes, see W. Geisler, 'Visual Perception and the Statistical Properties of Natural Scenes', *Annual Review of Psychology* 59 (2008), 167–192; A. D'Antona, J. Perry, and W. Geisler, 'Humans Make Efficient Use of Natural Image Statistics When Performing Spatial Interpolation', *op. cit.*; J. Burge and W. Geisler, 'Optimal Disparity Estimation in Natural Stereo Images' *op. cit.*. Most popular philosophical discussions of objective environmental priors cite such priors as that light is more likely to come from above than below. Such illustrations tend to obscure the variety and complexity of environmental statistics that are coded in perceptual systems. A huge amount of information is coded through evolution. This information is tuned through learning. Statistics that are likely to remain stable, such as statistical facts about motion governed by the principle of inertia, are likely to be innately coded. Statistics that vary during the lifetime of the organism, such as the distribution of prey, are likely to depend more on learning.

Measurement of some prior objective probabilities is much less difficult than that of others. Distance, being one dimensional, is relatively easy. Tilt, slant, binocular disparity values are harder but tractable. Hue is two dimensional. Size is two- or three-dimensional. 3-d shape is multi-dimensional. Higher-dimensional quantities are practically impossible to measure, because of the amount of data required to obtain reliable statistical measurements. Priors of any perceivable attribute are, however, in principle determinable, by way of dimensionality reduction—discussed in note 113. Similarly, for the idealization about initial retinal image registrations—which are extremely high-dimensional—cited in the next text paragraph.

I ignore surfaces of heavenly bodies, in referring to the horizon. I doubt that evolution provided for them in the same way that it provided for distances of surfaces on earth. Heavenly bodies can be seen. I doubt that human visual systems attribute definite distances to them.

Given the optical laws and given that the perceiver's eyes are open, the distribution of objective prior probabilities of distances of surfaces, together with a factor for degradation of the signal through noise, determines a distribution of probabilities that various initial registrations will occur for a surface at each given distance. From this joint distribution, a conditional probability distribution for the probability of any  $I$  for given a fixed  $D_j$  can be derived:  $p(I | D_j)$ .

So far, I have discussed objective probabilities relevant to Bayes's theorem. The probabilities are *objective* in that they do not concern the perspective of the perceiver, and involve no assignments of probability *in* any psychological system to a representational state.

Let us turn to interpreting Bayes's theorem as being about subjective probabilities assigned in a visual system. The probabilities are *subjective* in that they are assigned in a psychological system and concern the perceiver's perspective, in that they may not coincide with the objective probabilities, and in that some of them, at least, concern the probability of the occurrence or accuracy of a perceptual representational state.

I write ' $p^s$ ' for 'the subjective probability of'. Because  $I$  is not a representation with accuracy conditions, the meaning of ' $p^s(I)$ ' is not: the subjective probability of an accurate occurrence of  $I$ . ' $p^s(I)$ ' means: the subjective probability that  $I$  occurs, where  $I$  is any initial image registration possible in the system. The science presumes that visual systems can be modeled as establishing a probability distribution for occurrences of initial image registrations. Strictly, a visual system does not track full, complex initial image registrations themselves, but low-dimensional feature sets within those registrations that are relevant to the type of perceptual attribution at issue.<sup>113</sup>

Again, these are not assignments of probabilities by individuals, much less conscious assignments. The subjective probability is an anticipation setting or reliance setting. Different degrees of anticipation or reliance correlate with different subjective probabilities of an image registration's occurring. The anticipations are usually closely correlated with the objective probability of the registration's occurring (or a dimensionality-reduced aspect of its occurring).

Discounting noise, a perception as of a surface-distance is a computational function of the occurrence of a relevant initial image registration for a given individual at a given time. So the subjective prior probability of an *occurrence* of a relevant image registration will yield a subjective prior probability for an *occurrence* of the corresponding perception as of surface distance. Any subjective probability of

<sup>113</sup> Given the colossal number of possible retinal images, and because even local segments of images are very high-dimensional, image registrations may seem crazily unmanageable. The science presumes, however, that perceptual systems do manage. They are presumed to cut through the complexity of full image registrations via operations of *dimensionality reduction*. For an overview of such techniques, see Ali Ghosi, 'Dimensionality Reduction: A Short Tutorial', [http://www.math.uwaterloo.ca/~aghodsib/courses/f06stat890/readings/tutorial\\_stat890.pdf](http://www.math.uwaterloo.ca/~aghodsib/courses/f06stat890/readings/tutorial_stat890.pdf) (sourced August 2015). For illustrations of application of such techniques in vision science, see J. Burge and W. Geisler, 'Optimal Disparity Estimation in Natural Stereo-Images', *op. cit.*; J. Burge and W. Geisler, 'Optimal Speed Estimation in Natural Image Movies Predicts Human Performance', *Nature Communications* 6, Article number 7900 (2015). These articles show that only a relatively small, statistically manageable number of dimensions (on the order of 6–8) is needed to produce ideal solutions to several problems and to predict human performance—for such features as binocular disparity and speed.

the occurrence of a perceptual state—again discounting noise—will be close to and derivable from the subjective probability of the occurrence of the image registration that causes that type of state.

We are primarily concerned not with objective probabilities regarding distances, but with subjective probabilities concerning perceptual state types that attribute distances. I symbolize a perceptual state type as of the distance of a surface:  $D^*$ . Subjective probabilities apply to, among other things, perceptual (and belief) state types. A perceptual state that attributes a specific distance  $j$  is symbolized:  $D_j^*$ . There are many perceptual state types that attribute the same distance, most of which enter into different subjective probability distributions. To indicate a specific perceptual state type  $a$  that attributes distance  $j$ , I write  $D_j^{*a}$ . A specific attributive for distance will fix the attributed distance. So  $a$  determines  $j$ . I cite these items—a specific perceptual state type that attributes a distance and the distance itself—separately to emphasize their distinctness.<sup>114</sup>

I am interested in subjective probabilities for two features of perceptual state types. One is the probability of there being an *occurrence* of an instance of a perceptual state type, for a given individual at a given time. The other is the probability of an occurrence of a perceptual state type's *being approximately veridical* for a given individual at a given time.

I symbolize the subjective probability of there being an occurrence of (an instance of) perceptual type  $D^*$ :  $p^s(O(D^*))$ . I symbolize the subjective probability of the approximate veridicality of any occurrence of perceptual type  $D^*$ :  $p^s(V(D^*))$ . The subjective probability of the radical non-veridicality of any occurrence of  $D^*$  is  $[1 - p^s(V(D^*))]$ . All subjective probabilities are understood to be for a given individual at a given time.

How does a visual system develop prior probabilities for there being an occurrence of a given type of perception and for any such occurrence's being approximately veridical?

As I indicated, a visual system can derive the probability of an *occurrence* of an instance of the perceptual state type from the probability of an occurrence of the initial image registration that computationally determines occurrences of

<sup>114</sup> I discuss Bayesian models on the understanding that psychological laws, or law-like patterns, include specific perceptual state types. Different ways of representing a given distance have different psychological properties. They derive from different initial image registrations. They may be associated with different confidence levels. They may induce action at different rates of speed. They may have different priming or other associational histories. And so on. For more on this matter, see my 'Reply to Rescorla and Peacocke: Perceptual Content in Light of Perceptual Constancies and Biological Constraints', *Philosophy and Phenomenological Research* 88 (2014), 485–501. I think that these points are known and assumed in the science, by reflective scientists. However, even many Bayesian models in the science operate on an idealization—in effect, equivalence classes of psychological states. In our case, a relevant equivalence class would be visual perceptual states that represent a surface as being at a specific distance  $D_j^*$  (regardless of specific ways in which the distance is perceptually represented). Everything that I say about specific perceptual states, like  $D_j^{*a}$ , in what follows can be transposed to apply to coarse-grained equivalence classes of specific perceptual states and classes of particular initial image registrations. Strictly, the subjective probability assignments would be for averages of assignments to members of the equivalence classes, or paradigmatic assignments for the equivalence classes. The basic points that I will make apply under this transposition.



the perceptual state type. Strictly, if one averages across the variation of state types caused by the initial image registration together with the noise, one gets an expected, computed state type. (Noise causes different state-types given a type of initial registration.) I will continue to use shorthand. I will write of the perceptual state type computed from the initial registration, '*discounting for noise*'. The system can determine a probability for the occurrence of an initial image registration (or rather the relevant low-dimensional aspect of it) by tracking feedback frequencies.

A visual system builds estimates or assignments of probabilities for the approximate *accuracy* of an occurrence of a perceptual state type in a different way. Each possible initial image registration, or rather the relevant low-dimensional aspect it, has a distribution of subjective probabilities for the various distances. By applying a loss function and an estimator function, the system selects an attributive for one of these possible distances-*j*. (See note 108.) A loss function assigns penalties to various types and degrees of error. An estimator function operates on the loss function, according to some principle. The estimator function selects a perceptual attribution, taking into account the losses that would be incurred if the various respective attributions were inaccurate. The selection yields a perceptual state type *a* that attributes distance *j*. Given a probability distribution of distances, relative to an initial registration, and given an initial registration, the system can compute a probability distribution for strict accuracy and various degrees of error for instances of the perceptual state type that would be computed from that registration. The probability of approximate accuracy for instances of the perceptual state type is derived from this distribution. All and any instances of the state type receive the same subjective probability for being approximately accurate.

The subjective probability assignments for approximate accuracy derive from feedback about the type, sign, and degree of error. The feedback helps calibrate the perceptual system, compensating for error, especially in action. Given an evolutionary and individual history of feedback, the system has a wealth of information from which to compute the probabilities of various types and degrees of errors.<sup>115</sup>

I emphasize two points here.

First, the main shape of this distribution of prior probabilities is fixed by evolution, with adjustments in individual learning histories. This shape has a definite contour. Different types of errors are assigned different probabilities. Probabilistic relations between accuracy, different sorts of small errors, and larger errors are fairly well determined. Since the prior probabilities for perception are the main source of prior probabilities for perceptual beliefs, comparing prior probabilities for perceptual belief to those in a lottery case, or other cases in which one has no differentiating information, is off base. So the Keynesian distinction between uncertainty and risk is not central to understanding perceptual belief.

Second, the information used by the visual system to produce these *prior* probabilities suffices to produce the *conditional* probability of an approximately accurate

<sup>115</sup> There is evidence that human perceptual systems do build up subjective probabilities for veridicality, and the various types of possible error. See M. Ernst and M. Banks, 'Humans Integrate Visual and Haptic Information in a Statistically Optimal Fashion', *op. cit.*

perception of the relevant type given an initial image registration. Very likely, in the evolutionary pre-history of perceptual systems, the prior and conditional probabilities were established together from the same information. I will not focus on the genesis of these settings. I focus on Bayes's theorem as applying (a) to the already set-up probability distributions in a given individual at a given time, and (b) the use of those distributions when specific image registrations or perceptual states occur.

I have discussed symbolism and content of the prior subjective probabilities for the occurrence of an initial image registration and for the occurrence and veridicality of the relevant perception. We are preparing for an instantiation of Bayes's theorem. We now need to discuss the remaining term on the right side of the equation that constitutes the theorem.

$p^s(I | O(D^{a_j}))$  is the subjective probability distribution of the occurrence of an initial image registration  $I$ , given an occurrence of perceptual state  $D^{a_j}$ .  $p^s(I | V(D^{a_j}))$  is the subjective probability distribution for  $I$  given an approximately veridical occurrence of perceptual state  $D^{a_j}$ . When the perceptual type is fixed, a given  $I$ —say,  $I_r$ —is fixed. Recall that a state type  $a$  is an attribution of a specific distance  $j$ , in a specific way. So the subjective probability of  $I_r$ , given either  $O(D^{a_j})$  or  $V(D^{a_j})$ , will be near 1. For specific perceptual types derive only from corresponding specific initial registrations, (or from a class of initial registrations whose differences are psychologically irrelevant) for an individual at a time, discounting noise.

The left side of Bayes's equation is a subjective posterior conditional probability. We have been discussing two subjective probability assignments—one for occurrence, one for approximate veridicality.  $p^s(O(D^{a_j}) | I)$  is the subjective probability distribution of there being an occurrence of a perceptual state type  $a$  that attributes distance  $j$  in the specific way it does, given the occurrence of any initial image registration  $I$ . Again, for an appropriate registration  $I_r$ , this probability will be in the neighborhood of 1, since, discounting noise, a given perceptual state type will be computationally produced by a given type of initial image registration.  $p^s(V(D^{a_j}) | I)$  is the subjective probability distribution of the approximate veridicality of any occurrence of a perception  $D^{a_j}$  attributing distance  $j$  in the specific way  $a$ , given  $I$ .

$p^s(I | O(D^{a_j}))$  is the subjective probability distribution of the occurrence of an initial image registration  $I$ , given an occurrence of perceptual state  $D^{a_j}$ .  $p^s(I | V(D^{a_j}))$  is the subjective probability distribution for  $I$  given an approximately veridical occurrence of perceptual state  $D^{a_j}$ . When the perceptual type is fixed as a result of a computation, a given  $I$ —say,  $I_r$ —will be fixed, discounting noise. (Again, a state type  $a$  is an attribution of a specific distance  $j$ , in a specific way.) The system registers this relation.

So the subjective probability of  $I_r$ , given either  $O(D^{a_j})$  or  $V(D^{a_j})$ , will be in the neighborhood of 1. For specific perceptual types derive only from corresponding specific initial registrations, for a given individual at a given time, discounting noise. These conditional probability relations are set up in the system, at a given time, in advance of any given retinal image registration or any given perceptual state.

Given that  $I_r$  computationally leads to  $O(D^{a_j})$ , discounting noise, and given that the system can register this fact, the system will attach approximately the same subjective probabilities to  $I_r$  and  $O(D^{a_j})$ . Given either one of them, the subjective probability of the other will be close to 1, discounting noise.

The left side of Bayes's equation can be written:  $p^s(O(D^{*a}_j) | I)$ . This is the subjective posterior probability of an occurrence of a perceptual state type  $a$  that attributes distance  $j$  in the specific way  $a$  it does, given the occurrence of initial image registration  $I$ . Again, this probability will be in the neighborhood of 1 for specific image registration  $I_r$ . For, discounting noise, a given perceptual state type is computationally produced by a given type of initial image registration.

$p^s(V(D^{*a}_j) | I)$  is the subjective probability of the approximate veridicality of any occurrence of a perception  $D^{*a}_j$  attributing distance  $j$  in the specific way  $a$ , given  $I$ .

Certain instantiations of Bayes's theorem will be important to our assessments of the arguments that began this section. This is the joint subjective probability of  $O(D^{*a}_j)$  and  $V(D^{*a}_j)$ . It is written:  $p^s(O(D^{*a}_j), V(D^{*a}_j))$ . This is the probability of the conjunction: there is an occurrence of  $D^{*a}_j$  and that occurrence is approximately veridical. By the laws of conditional probability, this is equivalent to the joint probability of there being an occurrence of  $D^{*a}_j$  and  $D^{*a}_j$ 's being veridical, given that it occurs. This is written:  $p^s(O(D^{*a}_j), (V(D^{*a}_j) | O(D^{*a}_j)))$ . I am interested in this instance of Bayes's theorem

$$\frac{p^s\left(\left(O(D^{*a}_j), V(D^{*a}_j)\right) | I_r\right) = p^s(I_r | \left(O(D^{*a}_j), V(D^{*a}_j)\right)) p^s\left(O(D^{*a}_j), V(D^{*a}_j)\right)}{p^s(I_r)}$$

Let us put together two points already made. First, at any given time, the system attaches approximately the same subjective probabilities to  $I_r$  and  $O(D^{*a}_j)$ , where  $I_r$  is the initial image registration (or dimensionality-reduced aspect of it) that computationally produces the specific way of attributing distance  $j$ ,  $D^{*a}_j$ , discounting for noise.<sup>116</sup> Given either  $O(D^{*a}_j)$  or  $I_r$ , the system attaches approximately 1 to the other. Second, all subjective probabilities attach, in the first instance, to perceptual state types that attribute distance  $j$  in a specific way. So at any given time, the individual's system attaches the same subjective probability for approximate veridicality to all and any occurrent instantiations of a given perceptual state type, independently of whether there are occurrences of that type. These two points entail that the following four subjective probabilities are assigned approximately the same probability value, for a given individual at a given time:

- (a)  $p^s\left(V(D^{*a}_j) | O(D^{*a}_j)\right)$
- (b)  $p^s\left(V(D^{*a}_j) | I_r\right)$
- (c)  $p^s\left(\left(O(D^{*a}_j), V(D^{*a}_j)\right) | I_r\right)$
- (d)  $p^s\left(V(D^{*a}_j)\right)$ .

<sup>116</sup> I think that beliefs that derive from direct stimulation of the central visual areas of the brain are not of the same psychological type as visual perceptual beliefs. If this view were mistaken, the points of this paragraph could be rephrased so as to maintain their main thrust.

The significance of this fact is that the subjective probability for approximate veridicality of an instance of a perceptual state is not affected by occurrence of the type. A subjective probability value for approximate veridicality is coded in the system before any occurrence of the type.

It is useful to distinguish two aspects of subjective probabilities in dynamic operations. That is, there are two aspects to how subjective probabilities are affected from the time before an image registration occurs, or perceptual state occurs, and the time after it occurs.

The first is the way in which there is a change in the joint subjective probabilities for the conjunction of  $O(D^{*a}_j)$  and  $V(D^{*a}_j)$ , from the subjective *prior* probabilities for the conjunction, once there is an actual occurrence of  $D^{*a}_j$ —that is, given  $O(D^{*a}_j)$ . For simplicity, let us take ' $D^{*a}_j$ ' to stand for either the perception or the perceptual belief derived from it. Subjective prior probabilities for having an occurrence of a given type of perception and for having an occurrence of the associated perceptual belief will be almost the same. Similarly, the subjective probability for the veridicality of an occurrence of a perception of any given type, for a given attribution, will be almost the same as the subjective probability for the veridicality of an occurrence of a perceptual belief that attributes the same attribute and is derived from the perceptual attribution. They operate dynamically in the same way, absent background information.

The subjective prior probability that is a base for the change that we are interested in is the subjective prior probability for there being an approximately veridical occurrence of a specific state type— $p^s(O(D^{*a}_j), V(D^{*a}_j))$ . How is there a change in subjective probability, from this base, with the *occurrence* of a perceptual state or perceptual belief  $D^{*a}_j$ ? As the arguments that began this section claim, the subjective probability will be higher. For the occurrence of  $D^{*a}_j$  will be registered by the system. So the first “conjunct” will be established. Parallel considerations show, as the arguments claim, that the subjective probability for there being a radically non-veridical occurrence of  $D^{*a}_j$  will be higher than the prior probability. Low subjective prior probabilities are assigned both to there being an approximately veridical occurrence of the state type and to there being a radically non-veridical occurrence of the state type. These low probabilities derive largely from a low prior probability that an instance of the perceptual, or perceptual-belief, state type will occur at all. Given that one does occur, the subjective probabilities for there being an approximately veridical occurrence and for there being a radically non-veridical occurrence of the state type will both rise.

The second aspect of the dynamics of subjective probabilities concerns how subjective *conditional* probabilities operate before and after a perceptual state and perceptual-belief state occur. Subjective *conditional* probabilities for perceptual states and perceptual beliefs are just as much in place, at a time for an individual, as subjective *prior* probabilities are.<sup>117</sup> For Bayes's theorem to be of use, the system

<sup>117</sup> Again, both prior and conditional probability assignments are instilled through eons of interaction with the environment, with feedback on errors. The illusion (V) as of a hand that derives from one's being in sceptical scenario is not *specifically* allowed for in the low subjective probability of there being a radically false occurrence of an instance of the perceptual-belief type. Innate selection and learning are influenced

needs, before the occurrence of the next perceptual state, a model that derives both prior and conditional probabilities from the joint probability distribution.

What is the effect of an occurrence of an instance of a perceptual state or perceptual belief of the relevant kind on these antecedently established subjective conditional probabilities? The occurrence of the initial image registration causes an occurrent instantiation of the perceptual state and perceptual belief. The system registers these events. So  $p^s(I_r)$ , given  $I_r$  is 1 and  $p^s(O(D^{*a}_j))$  given  $O(D^{*a}_j)$  is 1. The subjective probability of  $I_r$  given an approximately veridical occurrence (or any occurrence) of  $D^{*a}_j$  is close to 1. So the subjective probability of the approximate veridicality of  $D^{*a}_j$  given a particular occurrence of  $D^{*a}_j$  is the same as the prior subjective probability of the approximate veridicality of any occurrence of  $D^{*a}_j$ . Thus (a), (b), and (c) have the same subjective probability value as (d). The system has already assigned subjective probabilities to the approximate veridicality and radical non-veridicality of all and any occurrences of given types of perceptual states and perceptual beliefs before they occur. The occurrences do not change those probability values. They do make the individual warranted in forming commitments to *de re* representational contents in accord with those values.

Assume that the subjective conditional probability (and subjective prior probability) for the approximate veridicality of  $D^{*a}_j$  is very high. Assume that the subjective conditional probability for radical non-veridicality is very low. Then the occurrence of the initial registration  $I_r$  and the occurrent instantiations of the perceptual state-type and perceptual belief-type  $D^{*a}_j$  leave the subjective probabilities for radical non-veridicality of these occurrent state instantiations well below the subjective probabilities for the approximate veridicality of the state occurrences. These probability values do not change with the occurrence of the perceptual state.

The high subjective posterior conditional probability of the approximate truth of any occurrence of a perceptual belief type, given an occurrence of the corresponding perceptual state and initial image registration, is set before the occurrence of the image registration. Equivalently, the low subjective posterior conditional probability of the radical falsehood of an occurrence of a perceptual belief, given an occurrence of the registration of the retinal image, is in place before the occurrence of the image registration. If the image registration occurs, the relevant perceptual states and perceptual beliefs occur. So given the probabilities in place before the occurrence of instances of the perceptual state types and perceptual belief types, the subjective probability of the approximate veridicality of any instances of those types will be high. The subjective probability of there being an occurrence of the perceptual belief type is higher after the occurrence than prior to the occurrence. So the subjective

only by feedback from errors that actually occurred in the molding of the system. Sceptical scenarios are not among such occurrences. So they are never among the *specific* radical falsehoods that a perceptual system allows for. A corollary of this point is that sceptical scenarios are not specifically within the hypothesis space. They are, however, species of the generic category of radically non-veridical representations (say, as of hands). As stressed in Section 2.3, sceptical scenarios are not relevant threats to entitlements. Rejecting ordinary radically non-veridical perceptual beliefs does not require a special reason, unless there are specific reasons to think that the radical falsehoods are threats. The sceptic must show that rejection of sceptical scenarios is different. I think that there is no reason to believe that, to be warranted in rejecting the sceptical scenarios, one needs a reason to reject them.

probabilities of there being an approximately veridical and of there being a radically non-veridical occurrence will both be higher after there is an occurrence of an instantiation of the perceptual belief type. It is true that the occurrence of the perceptual belief type is exactly what one would expect if one were to have a radically false occurrence of just that type. But all three subjective probabilities (for there being an occurrence of the perceptual type, for any occurrence's being approximately true, and for any occurrence's being radically false) are set independently of any occurrence. So the occurrence does not affect one's subjective probabilities specifically for approximate truth/radical falsehood. The perception-forming and perceptual-belief-forming systems ground one's entitlement to most perceptual beliefs as they occur.

Since the approximate truth of the perceptual belief is logically incompatible with its radical falsehood, the psychological system will place a much higher subjective probability on denials of the radical falsehood of occurrences of the perceptual belief than on acceptances of their radical falsehood. So the psychological systems' probability assignments, given that they are reliable, ground entitlement to reject an occurrence of the perceptual belief as radically false.

Let us return to the point that after occurrence of the image-registration, perceptual state, and perceptual belief, there *will* be a rise from the subjective unconditional probability of there being a radically false occurrence of the perceptual belief. This suggests that there will be some "confirmation" of the proposition that there is a radically false occurrence of the relevant type of perceptual belief. Confirmation here is an appropriate increase in subjective probability. It does not amount to prima facie warrant. Prima facie warrant marks a belief's propositional content as worthy of belief and as prima facie knowledge. Nothing so strong can be claimed for the increase in probability that there is a radically false occurrence of that type of perceptual belief.

Consider again:

(not-V) it is not the case that (V) I am in a sceptical scenario with a mistaken belief as of a body part and hand.

(not-V) is equivalent to:

(not-V') it is not the case that (V') I am having a belief as of a body part and hand and that belief is an illusion by virtue of my being in a sceptical scenario.

(not-V) is also equivalent to:

(DeMorganized not-V) Either I am not having a belief as of a body part and hand or I am having such a belief and it not an illusion by virtue of my being in a sceptical scenario.

These unapplied schemas do not themselves have truth-value. Hence they do not themselves have subjective probabilities for being true or false. The assigned probabilities apply to all and any instantiations of the schemas—all possible applications of them. In particular, they apply (at a time, for an individual) to all actual or possible applications of the progressive tense and applications within the perceptual belief content that<sub>1</sub> body part is a hand<sub>that2</sub>. Recall, yet again, that a visual perceptual belief of type that body part is a hand<sub>that</sub> is a very specific way of believing—with

specific body-part-, size-, and shape-attributions. (See note 114.) The schemas leave open whether there is an *occurrence* of the relevant belief-type. A high subjective unconditional prior probability in these cases is largely grounded in a low prior probability of there being an occurrence of a relevant perceptual belief type at all—radically non-veridical or not.

Suppose now that the relevant perception and perceptual belief do occur. There are *de re* applications of the propositional forms. Applications occur both in the perceptual belief and in self-attributions of it. Applications of tense in (not-V), (not-V'), and (DeMorganized not-V) occur. The first disjunct of (DeMorganized not-V) is falsified. Parallel points apply to (not-V) and (not-V'). The subjective probability of all three decreases. The subjective probabilities that had been schematically assigned to there being occurrences of that is a hand and I am under a brain in vat illusion in the belief that that body is a hand increase. In shorthand, when there is an applied occurrence of the perceptual-state type and perceptual-belief type, the subjective probabilities of all three conjunctions [occurrence and approximately veridical], [occurrence and not radically non-veridical] and [occurrence and radically non-veridical] all go up. They go up because the psychological system has information that shows that the first conjunct is true. These points accord with the claim in step (2<sub>B</sub>) of the first argument. The occurrences of the relevant perceptual-state type and perceptual-belief type are exactly what one would expect if one were to have a radically non-veridical perceptual belief of the relevant type (even a brain-in-vat-illusion).

But these changes in subjective probability do not bear in any way on the subjective probabilities already assigned to the approximate truth and radical falsehood of any application of the schematic forms that body part is a hand<sub>that</sub> or I am under a brain in vat illusion in the belief that that body part is a hand<sub>that</sub>. The changes in probability so far discussed derive entirely from the belief's propositional content's coming into existence. The changes in subjective probability that the arguments fix upon are not changes that focus on the subjective probability of the approximate truth or radical falsehood of any occurrent perceptual belief. All perceptual beliefs are occurrent.

As I remarked, there are subjective probability assignments that do assess the approximate truth and radical falsehood of any occurrence of an instantiation of the perceptual belief type. These assignments are already in place before the occurrence of the perceptual belief. When the applied propositional content that instantiates the perceptual belief type *does* occur, those subjective probability assignments for approximate truth and radical falsehood attach to the occurrence. They attach to that<sub>1</sub> body part is a hand<sub>that2</sub> and to I am<sub>0</sub> under a brain-in-vat illusion in the belief that<sub>1</sub> body part is a hand<sub>that2</sub>.

One such assignment is the assignment that is *conditional* on the occurrence of a relevant image registration, or the occurrence of an instance of a relevant perceptual state type. Again, such a *conditional* assignment is in place in the psychological system before the occurrence of the *de re* applied perceptual state and perceptual belief. When the condition is met—when there is an occurrent instantiation of the perceptual state type—the relevant degree of subjective probability attaches to the occurrent perceptual belief content. There is also a subjective *non-conditional* probability assignment in the psychological system before the occurrence of the relevant

image registration, perceptual state, or perceptual belief. This assignment is for the approximate truth (and correspondingly, the radical falsehood) of all and any occurrences of the relevant perceptual belief type. This assignment has approximately the same value as that of the conditional probability assignment. For, roughly, discounting for noise and pathology, the perceptual belief content occurs if and only if the perceptual state type occurs. Conditional and unconditional assignments are typically very high for approximate truth and very low for radical falsehood.

The assignments to occurrent, *de re* applied perceptual beliefs yield entitlement to the perceptual belief, and to rejection of its being radically false. The assignments (contingently but non-accidentally) track the objectively reliable approximate truth of instances of the perceptual belief types. They provide confirmation of the perceptual belief's representational content when it occurs. Being assigned in the same psychological system at the same time, they limit the slight rise in the probability of a radically false occurrence that derives simply from there being an occurrence of the relevant type. Although the levels of subjective probability are fixed in advance and independently of any actual occurrence, an individual receives warrant for believing the belief's content only with its occurrence. An occurrent, *de re* applied perceptual belief content that can be approximately true or radically false does not exist before, or independently of, its occurrence. Lacking at least an anticipation of a specific occurrence (a case that I shall return to), no occurrent application can be individuated before its occurrence. There is no full-fledged perceptual belief content to be committed to, or for a subjective probability to be assigned to.

Again, at any given time, the levels of subjective probabilities for the approximate veridicality and radical non-veridicality of any and every occurrence of a perceptual state and perceptual belief are assigned in advance of occurrence. In this sense, the assignments are not only unconditional, but "prior". It is not a subjective prior probability in the sense that one already has some degree of belief or disbelief for a specific propositional content in advance of the occurrence. For the content of the belief does not exist before its occurrence.<sup>118</sup>

The subjective probability assigned to an occurrent, *de re* applied perceptual belief can be changed by acquisition of new information beyond that encoded in the perceptual belief. It is not changed by the occurrence of the perception or perceptual belief.

As noted earlier, the subjective probabilities for the approximate truth—and for the radical falsehood—are *empirical*. They apply to any occurrence of a perceptual state or perceptual belief of a given type *before*—and independently of whether—the state type is instantiated. They are empirical because they depend for the force of the warrant that they help establish on past sensory interactions between perceptual system and environment, partly in the evolution of the perceptual system. (See note 110.) The subjective probabilities are set in the perceptual—and perceptual-belief-forming systems, not from the armchair.

<sup>118</sup> The occurrent *de re* applied perceptual beliefs are flat-out approximately true or radically false, once they do occur. Their truth or falsity is not relative to a time, or to anything else.



These points ground assessment of the arguments set out at the beginning of the section.

I begin with the intuitive argument (1<sub>B</sub>)-(5<sub>B</sub>). It is important, in assessing steps (2<sub>B</sub>) and (3<sub>B</sub>), that one be clear exactly what they mean. (not-V) must be interpreted with some care. If (not-V) is regarded as a schema that leaves open whether there is a *de re* applied perceptual state and perceptual belief as of a hand, then as the argument claims, there is a decrease in subjective probability when (PB)-(PH) and (M) occur. However, the schematic interpretation of (not-V) is not the one intended in the Moorean argument (M)-(A)-(not-V). In the Moorean argument, (A) and (not-V) are interpreted as *de re* applied in a context. They are interpreted:

- (A<sub>1</sub>) that<sub>1</sub> body part is a hand<sub>that2</sub> only if it is not the case that the current belief that<sub>1</sub> body part is a hand<sub>that2</sub> is a brain-in-vat illusion

and

- (not-V<sub>1</sub>) it is not the case that (V') the current belief that<sub>1</sub> body part is a hand<sub>that2</sub> is a brain-in-vat illusion.

These propositional contents (not schematic contents) occur only because there are occurrences of (PB)-(PH) and (M). The subjective probability of the truth of (not-V) so interpreted does not decrease. Although neither (not-V<sub>1</sub>) nor (M) existed before their occurrence, a given subjective probability was set in advance to be assigned to them. For there is a subjective probability assigned to the approximate truth, and correspondingly the radical falsehood, of every and any occurrence of the perceptual belief type that body part is a hand<sub>that</sub> for any applications of the thats. Indeed, since body part and hand are not scalar attributes, one can perhaps carry out the reasoning on truth and falsehood (not their approximations) *simpliciter*. The relevant subjective probability does not change with the occurrence of applications. So, given the way that (not-V) is interpreted by the Moorean argument—as (not-V<sub>1</sub>)—the subjective probability assigned to the truth of (not-V) does not change with the occurrence of (PB)-(PH) or (M).

Prior to its occurrence, strictly speaking, one had no entitlement to believe (M). One could not have an attitude toward (M) or (not-V<sub>1</sub>) themselves, since the propositional contents, which include the occurrent applications, did not exist. With (M)'s occurrence, one's subjective probabilities that a true and that a false occurrence *occur* do increase. At the same time, the antecedent, high subjective probability for the *truth* of any occurrence of (M)—and low subjective probability for the *falsehood* of any occurrence of (M)—kick in. Those assignments place truth of the occurrence of (M) well above .5 and the falsehood of the occurrence well below .5. One gains an entitlement to believe (M). One gains confirmation, but not against a preceding baseline of a different assignment of subjective probability. One gains prima facie default entitlement *in* the occurrence of the perceptual belief.

What of the reasoning from (M) and (A<sub>1</sub>) to (not-V<sub>1</sub>)? One can use the entitlement to (M) and the reason to believe (A<sub>1</sub>)—an obvious logical or semantical truth—to provide a reason for (not-V<sub>1</sub>). Again, the subjective probability of (not-V<sub>1</sub>) does not

strictly rise against a baseline that preceded the occurrence of (M) or (not- $V_1$ ). (not- $V_1$ ) itself did not exist to have a probability. The argument does provide a transmission of warrant from (M) and ( $A_1$ ) to (not- $V_1$ ).<sup>119</sup>

To simplify all this, let us walk through the intuitive argument (1<sub>B</sub>)-(5<sub>B</sub>) to bring out again the intuitive points about where it goes wrong.

Step (1<sub>B</sub>) If I gain warrant for believing a proposition, my confidence in the proposition should increase

does not apply straightforwardly to entitlements to perceptual belief. I gain warrant to believe (M) That<sub>1</sub> body part is a hand<sub>that2</sub> only with the occurrence of the belief. Usually, one has no level of confidence in (M), the occurrent *de re* belief itself, before its occurrence. In that sense, there is no rise in confidence. There is, however, a level of confidence in the truth of any instance of the schematic belief-type, That body part is a hand<sub>that</sub>. That level is set before the occurrence of any instance of that schematic type. So with an occurrence, that level does not rise. What rises is the level of my antecedent confidence in there being an occurrence of an applied instance of the schematic belief-type That body part is a hand<sub>that</sub>. It is true that I had very low confidence of there being any occurrence of the schema before (M) occurred, but higher confidence afterwards. The belief that<sub>1</sub> body part is a hand<sub>that2</sub> is essentially an occurrent belief. It is also true that I had very low confidence in there being an occurrence of the negation of (M) before its occurrence and higher confidence afterwards. But the relative levels of confidence, regarding veridicality, in (M) and in its negation, are fixed in advance, via assignments to any occurrence of the relevant schematic types. The confidence level in (M) is much higher than the confidence level in its negation.

Step (2<sub>B</sub>) If (M) and (A) provide warrant for (not-V), my confidence in (not-V) should increase

is problematic in the same way. My confidence in the truth of any instance of the schema (not-V) (it is not the case that (V) I am under a brain-in-vat [belief] illusion as of a body part and hand) is set in advance. That level of confidence is vastly higher than the level set in advance for any occurrence of the schema (V). Those levels do not change with the occurrences of the relevant perceptual state and perceptual belief. The Moorean argument hinges on confidence in the truth of an actual occurrence of (not-V). I gain warrant for believing (not-V) only when it comes

<sup>119</sup> One can also reason from one's entitlement to reject the falsehood of (M) back to accepting the truth of (M). For the perceptual-belief-forming system assigns a very low subjective probability to the falsehood of (M). So one is entitled to believe: (M) is not false. One has reason to believe: (M) is not false if and only if (M) is true. (Clearly, if needed, one could substitute 'not true' for 'false' in the argument.) One has reason to believe: (M) is true if and only if that<sub>1</sub> body part is a hand<sub>that2</sub>. (These latter two beliefs are obvious logical or semantical truths.) So one has reason to believe: that<sub>1</sub> body part is a hand<sub>that2</sub>. This is warranted reasoning that supports its conclusion. It is obviously less natural than the Moorean argument from (M) and ( $A_1$ ) to (not- $V_1$ ). It is less natural because it starts not with the ground floor perceptual belief, but with a meta-belief about (M). I mention the argument to stress that the perceptual and perceptual-belief-forming systems simultaneously assign subjective probabilities to all and any instances of a perceptual-belief-type for a whole distribution of cases. For non-scalar properties like *body part* and *hand*, I assume that the argument need not appeal to approximate truth and radical falsehood.

into existence—only with its occurrence. So strictly speaking, Step (2<sub>B</sub>) is mistaken. The level of confidence for any occurrence of (not-V) is set in advance of, and independently of, its occurrence, and hence in advance of my having warrant to believe that specific occurrent propositional content.

- (3<sub>B</sub>) Given the sources of entitlement for (M)—the perceptual states,  
 (P-B) that<sub>1</sub> body part [with an attendant relevant shape, also perceptually referred to]  
 (P-H) that<sub>2</sub> hand—  
 my confidence in (not-V) should decrease

is problematic in the same way. It is true that my confidence in there being an *occurrence* of a perceptual state and perceptual belief predicted by (V) goes up. But my confidence in the truth of any instance of (not-V) is set in advance of its instantiation or occurrence, hence in advance of my having warrant to believe that specific occurrent propositional content. Of course, occurrences of perceptions can alter subjective probabilities set for future occurrences. Such changes include perceptual learning and learning associated with perceptual belief. Usually, significant, stable changes in subjective probabilities, for perceptions and perceptual beliefs (understood in my strict way), are slow to take hold.

Since Steps (1<sub>B</sub>)-(3<sub>B</sub>) produce no reason to be puzzled by entitlements to perceptual belief, conclusions (4<sub>B</sub>) and (5<sub>B</sub>) of the argument cannot be drawn. In particular, contrary to step (5<sub>B</sub>), one does not need independent warrant, beyond that provided by a reliable, well-functioning perceptual-belief-forming system, to believe (M)—That<sub>1</sub> body part is a hand<sub>that2</sub>. Insofar as the subjective probabilities of those systems track objective probabilities, as they have been shown in the relevant science to do, those systems themselves help generate warrant for perceptual beliefs.

These remarks about argument (1<sub>B</sub>)-(5<sub>B</sub>) carry over to the first of the two Bayesian arguments set out at the beginning of this section. What of the second Bayesian argument?

This argument suffers from difficulties that undermined (1<sub>B</sub>)-(5<sub>B</sub>) and the Bayesian counterpart of (1<sub>B</sub>)-(5<sub>B</sub>).<sup>120</sup> The argument is correct, strictly speaking, in claiming that ‘the probability of (M) given (PB)-(PH) is less than the probability of (not-V)’. For (not-V)—interpreted as (not-V<sub>1</sub>)—can be made true not only by (M) but by any illusion besides the brain-in-vat illusion. The argument errs, however, in its claim, ‘So (PB)-(PH) can make one warrantably confident that (M) only if one is already warrantably confident that (not-V)’. One has a warrant to be confident in (M)

<sup>120</sup> One problematic claim is: ‘And the probability of (not-V) given (PB)-(PH) is less than the probability of (not-V)’. If (not-V) is interpreted as the Moorean argument interprets it—as (not-V<sub>1</sub>)—then the probability of the truth of (not-V<sub>1</sub>) and the probability of the truth of [(not-V<sub>1</sub>) given (PB)-(PH)] are almost the same. For the current perceptual belief that<sub>1</sub> body part is a hand<sub>that2</sub>, mentioned in (not-V<sub>1</sub>) does not normally occur without (PB)-(PH), and the system is sensitive to this fact. The ‘normally’ is, of course, a hedge. (not-V<sub>1</sub>) could have occurred without (PB)-(PH), if it were caused by non-standard perceptual states together with noise. So strictly, the argument’s claim here is correct, but not for the reasons it envisaged. The difference in probability that the claim is correct about can be factored in, in advance of the occurrence of the perceptual states.

through the occurrence of (PB)-(PH) and (M). Normally, one cannot be warrantably confident in (not- $V_1$ ) before the occurrence of (M) or (not- $V_1$ ). Those contents do not exist before their occurrences. On the other hand, the subjective probability values to (M) and (not- $V_1$ ) are assigned to any occurrences of the propositional schemas in advance. So the relevant subjective probabilities for (not- $V_1$ ), as well as (M), are fixed independently of the occurrence of (M). The argument's conclusion—

- (5<sub>B</sub>) One needs *independent* ground to believe (not- $V$ ) in order to be warranted in believing (M) through having perceptions (PB)-(PH)

—is, however, mistaken. The subjective probabilities for (not- $V_1$ ) and (M) are fixed independently of the occurrence of (M). They are not fixed independently of one another. The perceptual-belief-forming system carries probabilities for strict truth, small errors, and radical errors all together. One can reason from (M)'s being true to its not being false, and from its not being false to its not being false in the particular way envisaged by (not- $V_1$ ).<sup>121</sup> The force of the reasoning does not depend on any propositional support, beyond the entitlement to (M).

I have remarked earlier (note 69) that my account of entitlement is not a dogmatist account. I want to expand the point. Pryor's explication of dogmatism is as follows:

The dogmatist about perceptual justification says that when it perceptually seems to you as if  $p$  is the case, you have a kind of justification for believing  $p$  that does not presuppose or rest on your justification for anything else, which could be cited in an argument (even an ampliative argument) for  $p$ . To have this justification for believing  $p$ , you need only have an experience that represents  $p$  as being the case. (See note 69.)

<sup>121</sup> White claimed that the effect of his arguments was to require that one must have independent ground to believe that one is not having a brain-in-vat illusion. He claimed that the ground had to be independent of the occurrence of the perceptual state. Roger White, 'Problems for Dogmatism', *op. cit.*, 533–534: 'So its appearing to me that this is [a] hand can render me justifiably confident that it is a hand, only if I am already confident that it is not a fake-hand.' White cites his point as supporting the view of C. Wright, '(Anti-) Sceptics Simple and Subtle: G. E. Moore and J. McDowell', *op. cit.*. White took his claim to show that one needs an independent *reason*, a warranted propositional basis, to accept (or apply a high probability to) the negation of a claim that the perception is radically inaccurate and the belief is radically false. In fact, the relevant probabilities, though independent of the actual occurrence of the perceptual state and perceptual belief, are assigned in the perceptual system and perceptual-belief-forming system themselves. They are empirical, even fundamentally perceptual. The individual need not be able to think the assignments. They are part of the entitlement system, not part of a reason-based, justificational system, or a system of evidence citation.

Contrary to what seems to be suggested in the quote, the confidence ratings of the [at least approximate] truth of the perceptual belief and of the denial of radical error in the perceptual belief are set together. Epistemically, neither the assignment of a high subjective probability to the approximate truth of the belief nor the assignment of a high subjective probability to the belief's not being radically false is prior to the other. Neither occurs before the other. The subjective probability assignments that favor approximate truth and disfavor radical falsehood are set independently of, and in advance of, the actual occurrence of the perceptions and beliefs. But both are set within the perceptual system and perceptual-belief-forming system themselves. White and Wright miss this fact. The relevant subjective probability levels are independent of and prior to the relevant *de re* applied beliefs' and perceptual states' occurrences. They are independent of its actually appearing to one that there is a hand. They yield entitlement only with the occurrences. But they are set within the perceptual system and perceptual-belief-forming systems themselves in advance of occurrences. They are not independent in the sense that they are background beliefs, supplementary to what goes on in the perceptual or perceptual-belief-forming systems.

I agree that one normally has a warrant for a perceptual belief if one has a perceptual state that represents, non-propositionally, what the propositional content of the belief represents. I also believe that the warrant does not presuppose or rest on one's warrant for anything else that could be cited in an argument for that propositional content. What I reject is that the warrant requires nothing other than having the corresponding perceptual state or experience. A perceptual state does not *in itself* provide any warrant at all for a perceptual belief. An entitlement requires something more than having the corresponding perceptual state.

The distinction between the distinctive claim of dogmatism and just the view that, typically, warrant for perceptual belief does not depend for its force on reasons or propositional background protection is widely elided in the literature. Pryor does not blur the distinction.

The occurrence of the perceptual state is not *in itself* the source of warrant. Again, in itself, the occurrence in itself provides no warrant at all for the belief. The perceptual state must derive from a well-functioning, reliable perception-forming competence, whose reliability is grounded in the environmental conditions in which the content-nature of the competence is determined. Reliability is not determined by the perceptual state itself. (See Section 2.2.) The same type of perceptual state metaphysically could be reliable or not in different content-constituting conditions. In some possible content-constituting (or content-sustaining) conditions, for a given state, practical evolutionary forces metaphysically could have engendered the same representational content, but yielded a different degree of reliability as regards veridicality. Reliability is determined by deeply non-accidental but metaphysically contingent connections of the system of production to its actual content-constituting conditions.

Individuals rely, in actions and responses, on perceptual states' and perceptual beliefs' being veridical. Their psychologies have grades of reliance on or anticipation of the veridicality of these states and beliefs. The numerical subjective probabilities for veridicality postulated by Bayesian models are just numbers put on these grades of reliance or anticipation.

Subjective probabilities derive from feedback over eons of evolution. If grades of reliance correspond well to objective probabilities of being right, and those objective probabilities are high, relevant perceptual beliefs are typically warranted. Connections between reliance and objective reliability, in a well-functioning system for producing perceptual states and perceptual beliefs, are the keys to entitlement—not the perceptual states in themselves.

I doubt that, in holding that the occurrence of a perceptual experience *in itself* warrants the belief, a dogmatist can correctly resist the arguments that began this section. At any rate, none of the defenses against them that I have seen seem to me to work. It is important to the present counter-argument, and to understanding the nature of warrant for perceptual belief, that warrant derives from empirical subjective probability assignments that are in place independently of the *occurrence* of relevant perceptual beliefs. Those assignments are consequences of eons of feedback regarding results (mainly in action and reaction) of sensory interactions with the environment. Dogmatism was right to hold that perceptual beliefs do not need help from

further actual or potential beliefs. It was wrong to hold that perceptual beliefs are warranted by nothing more than occurrences of perceptual states (or experiences).

White was right to reject dogmatism. That is, he was right to hold that the warrant for perceptual belief does not reside in the perceptual state's occurrence in itself. He was wrong to apply the criticism to my entitlement account. White, following Wright, thought that support, beyond occurrence of the perception in itself, had to come from *background, propositional* reasons available to the believer. This view has never been plausible. What I have tried to do is to show in detail how warrant comes from aspects of the perceptual system and perceptual-belief-forming system, and their relations to the environment. The warrant is not present unless one has a perception. But the connections that give the warrant force are independent of whether the perceptual state occurs. They are in place prior to its occurrence.

I briefly discuss two issues associated with the account that I have given.

The first concerns the point that usually the representational content of a perceptual belief does not exist before the occurrence of the belief. I want to discuss this point in relation to the subjective prior probabilities assigned by individuals on the basis of considerations that antedate the occurrence of a perceptual state. These are the assignments that are normally discussed in the literature on subjective probability and confirmation.

The representational contents of specific perceptual beliefs are individuateable only if their occurrent applications are individuateable. An individual or psychological system usually cannot individuate applications in advance of their occurring. So subjective probability assignments typically do not apply to specific propositional contents for perceptual beliefs until they occur. In this sense, there is no subjective prior probability for the approximate truth or radical falsehood of the *specific* occurrent propositional content—prior to its occurring. This point is rendered less momentous by the point that the subjective prior probabilities for approximate veridicality and radical non-veridicality are set for *any* and every propositional content of the relevant type, independently of whether it occurs. The individual gains entitlement to accept the propositional content of a specific perceptual belief only when the full content, including applications, occurs, even though the subjective probability that underlies the entitlement is set before the occurrence.

Let us consider two cases in light of these points.

First, suppose that I know somehow that 80% of the spherical bodies in a large bowl of such bodies are red. I know that the others are green. Suppose that I perceive, by touch, one of the bodies, but do not have a visual perception of its color. Then I can assign a conscious prior subjective probability of .8 to the veridicality of the propositional content:

(\*) This<sub>1</sub> body is red.

After one looks at the body and has a perception as of its being red, one assigns a probability for the veridicality of (\*) that is considerably higher than .8. The conditional probability of (\*) given a visual perception as of the body's being red is presumably much higher than .8. For the probability of any occurrence of the perceptual state's (and the visual perceptual belief's) being approximately veridical

is much higher than .8, and one can register the occurrence of the perceptual state and use this registration to provide confirmation for (\*).

The intuitions set out by the arguments at the beginning of the section seem to apply. The subjective probability for the veridicality of (\*) goes up with the occurrence of the relevant perception. What of the probability of there being a radical illusion as of a red body? The subjective probability for the veridicality of

(\*illusion) I have a radically illusory perceptual state as of a red body, when I look at that<sub>1</sub> body

also goes up, relative to the prior for the schematic version of (\*illusion), when the relevant perception as of a red body occurs. For (\*illusion) could be falsified both by a perception (veridical or not) as of a green body and by a veridical perception as of a red body.

The points already made reapply. There is a prior assignment in the visual system of a very high probability to the approximate veridicality of any occurrence of the relevant perception as of a red body. There is a prior assignment of a very low probability to the radically non-veridicality of any occurrence of the perception. So the rise in subjective probability from the individual's conscious prior probability assignment to (\*illusion) is limited to a level vastly below .5. So (\*illusion) is not confirmed, in the sense of warranted. In fact, after the perception, one has a warrant to reject (\*illusion).

It is perhaps worth noting that (\*) is not the content of a perceptual belief. A perceptual belief takes only concepts that are conceptualizations of perceptual attributives, and it takes over all demonstrative-referential applications that are perceptual applications of perceptual attributives that it conceptualizes and incorporates analogs of those perceptual applications in the perceptual belief. (\*) does not meet these conditions. (\*) is not a possible perceptual-belief content because red is not applied demonstratively in the way that it is in a perceptual belief (and in the underlying perception). That is clear from the example. The representational content does not function to refer perceptually to an instance of red in a *de re* perceptual application. So (\*) does not have the logical form of a perceptual belief. The logical form must have a demonstrative application that red guides. The probability is assigned so as to presuppose that no instance of red is perceived. Hence there is no perceptual demonstrative-like referential application. (\*) is confirmed and warranted not just by a perception as of red, but by a perceptual belief that<sub>1</sub> body is red<sub>that2</sub> or by some other empirical belief that traces back to a perceptual belief.

I said that '*typically*' propositional contents of perceptual beliefs cannot be individuated in advance of their occurring. There are probably exceptions.<sup>122</sup> Let us consider a second case. Suppose again that one feels a spherical body as a spherical body, but does not see it. One might form a belief about the body's color and have a subjective probability for the attribution of the color. Such a belief will likely lack the representational content of a perceptual belief (as with (\*)). For it will likely lack the visual attributives, and applications to instances of the color, that

<sup>122</sup> I owe the main outlines of this case to Michael Rescorla.

would occur in a visual perceptual belief. Still, one might form an anticipatory belief that does have the same representational content as an anticipated visual perceptual belief. To have the same content, the anticipatory belief would have to involve only conceptualizations of visual perceptual attributives. And it would have to have the same referential applications as the visual perceptual belief. These would have to depend for their identity and reference on applications in the anticipated perception whose attributives are conceptualized in the belief. Perhaps one could visually imagine the anticipated body, shape, and color shade and bring up the same visual conceptual attributives that later occur in conceptualizations in the content of the visual perceptual belief. One could anticipate the events of *de re* application to the instances of body, sphericity, and color shade. Perhaps one could individuate the perceptual-belief applications before they occur by knowingly anticipating their occurrence, and by putting oneself in circumstances to cause their occurrence. A subjective probability could then be assigned to this content. The content is that of an anticipation of a perception belief, not yet that of a perceptual belief. The basis for the probability assignment might be in memory, interlocution, or background information. Then one has a subjective prior probability that attaches to that very content, before the perceptual belief is formed. Such a case is rare, but I conjecture possible.

The account that I have given applies to this case. Suppose that one has the anticipated visual perception as of the color shade. Suppose that one defeasibly forms the corresponding perceptual belief. Then the subjective probability of there being an occurrence of an approximately true attribution of the color shade will go up. The subjective probability of there being an occurrence of a radically false attribution of the color shade will also go up.

The subjective probability of the approximate truth of any occurrence of the propositional content—and correspondingly, the probability of the radical falsehood of any occurrence of that content—that is contributed by the perceptual-belief-forming system are fixed in advance.

Assume, that visual perception *prima facie* outweighs memory, interlocution, and other background information. The *prima facie* warrant for the approximate truth of the visual perceptual belief's content outweighs the warrants for the content when the perception is just anticipated. And the subjective probability assigned to the content after the occurrence of the perception is higher than the assignment to the content when the perception was just anticipated. The subjective probability for the content's approximate truth therefore increases with the occurrence of the visual perception. Since the subjective probability for any occurrence of the anticipated content's radical falsehood is capped well below .5, one is warranted after the visual perception not only in believing in the approximate truth of the visual color shade attribution, but also in denying its radical falsehood. If, in the abnormal case in which memory, interlocution, or other background information outweigh perceptual belief, the subjective probability of the representational content's approximate truth does not change with its vision-based occurrence.

The second issue concerns the relation between subjective probability assignments that occur unconsciously in the perceptual or belief-forming systems, and



assignments that are attributable to the individual. The latter are commonly elicited by testing how an individual is inclined to bet. Some might worry that unconscious, possibly “sub-personal” (or, to allow for animals, “sub-individual”), assignments are not relevant to epistemology.

Probability assignments to occurrences of non-representational sensory *registrations* are sub-individual. Veridicality assignments to perceptual states and perceptual beliefs are largely unconscious and not deliberately controlled. The anticipations and reliances that the subjective probabilities mark are broadly like an individual’s willingness to bet. Assignments to perceptual states posited in perceptual psychology are more ingrained, perhaps less accessible to conscious estimates, than Bayesian epistemologists tend to represent subjective probability assignments as being. Perceptual psychology’s use of such assignments is to that extent more realistic.

When and whether unconscious subjective probability assignments to perceptual states or perceptual beliefs are ascribable to individuals or only to their sub-systems requires further clarification of the individual/sub-individual distinction. I think that the account that I have given does not hinge on how to draw and apply the distinction.

Assignments set in psychological sub-systems are subject to normative standards. The degree of reliance on specific perceptual beliefs depends primarily on the degree of reliance on corresponding perceptual states. Degrees of reliance on these states for their approximate veridicality appear to be nearly optimal in a range of central cases. (See notes 54, 104.) The assignments to perceptual states are, in those cases, very close to what, from an epistemic perspective, they should be, in normal conditions. Deviations from strict veridicality mostly result from physical limitations of acuity and noise. Individuals who rely for the strength of their perceptual beliefs on the specific natures of their perceptual states are likely to have degrees of confidence that fit well with objective conditions.

Normative epistemology applies to psychological capacities. Discounting the subjective probability assignments that I have discussed would leave epistemology inapplicable to actual psychologies. The arguments that began this section are ingenious. But they illustrate the hazards of doing the epistemology of perceptual belief without reflecting on the psychology of perception. The argument and previous replies to it are unrelated to applications of subjective probability theory in the relevant domain. Epistemology must be better informed by psychology.

## 2.6 Conclusion

Entitlement to perceptual beliefs is the basis for all empirical epistemic warrant. This simple statement is part of a large shift that has been occurring over the last several decades away from much twentieth-century philosophical work on mind and warrant. The shift displaces rational capacities, and consciously accessible aspects of psychologies, from being nearly the sole focus of a philosophical account of mind and knowledge to being just two important factors. The shift forces conscious rational capacities to share focus with largely unconscious psychological sub-systems and

their relations to a wider non-psychological environment. Similar displacements have figured in the earlier histories of science and philosophy.<sup>123</sup>

The shift regarding mind, embodied in anti-individualism, is a return to an older position, initiated by Aristotle and dominant in philosophy until the twentieth century. This view is that the natures of most psychological states are constitutively determined by relations that the individual or the individual's psychological systems bear to a wider environment. Through causal-interactive patterns, some of which reach over evolutionary time, physical subject matters stamp themselves into the representational natures of psychological systems and psychological states.<sup>124</sup> Development of the view and good arguments for it were not prominent in the tradition. But the view itself was a majority position in philosophy, until the twentieth century.

The shift regarding warrant—and its ancestor, following rules for obtaining scientific cognition—began in late-mid-twentieth century.<sup>125</sup> The idea that epistemic warrants for an individual are not entirely determined by the natures of and relations among an individual's psychological states is not a return to a dominant view. At least since the early modern period, epistemic internalism dominated epistemology.

To specify the displacements: Anti-individualism shows that representational mind is constitutively embedded in a wider world. The physical environment has a lead role in making perceptual states and empirical beliefs what they are. It does so partly through patterns of causal interactions with sub-representational sensors. These patterns help type-individuate the representational states. They also give such states their subject matters. The environment has a further lead role in yielding epistemic warrant for perceptual beliefs. It contributes in two ways. One is to set the conditions in which reliable veridicality counts toward warrant. The other is to help set a knowledge-inducing and warrant-inducing route back to the subject matter of perception and perceptual belief. So environment-mind connections that constitutively helped determine the representational natures of perceptual states and

<sup>123</sup> Obviously, the Copernican, Darwinian, and Einsteinian revolutions all, in various ways, displaced humans and ordinary human experience from a domineering role in determining truths about a subject matter. Parallel changes in philosophy—the separation of philosophy from religion, the recognition of the brute non-rational character of the physical world, the acknowledgment that philosophy is not the queen of the sciences, and the decline of idealism—participate in this large current of displacement. Ironically, Kant appealed to the Copernican revolution as an analog of his regressive defense of a form of idealism about the scientifically cognized world. However, Kant made major contributions to the other three shifts in philosophy just cited.

<sup>124</sup> See *Foundations of Mind*, *op. cit.*, essays 4–9 and *Origins of Objectivity*, *op. cit.*, chapters 1 and 3. Anti-individualism applies, in a different way, to psychological states about mathematics and logic.

<sup>125</sup> See the authors cited in note 70. W. V. Quine urged that epistemology use scientific knowledge regarding perception and discovery. See, for example, his 'Epistemology Naturalized', in *Ontological Relativity* (New York: Columbia University Press, 1969). He saw himself as suggesting that science replace traditional epistemology. I think that this view was deeply misconceived. He reflected too little on the difference between epistemic norms and the topics of ordinary empirical scientific knowledge. And he did not understand the relevant psychology. Still, he contributed by advocating connecting epistemology with science. After Quine, Alvin Goldman in 'What is Justified Belief', *op. cit.* and *Epistemology and Cognition*, *op. cit.* and Fred Dretske, *Knowledge and the Flow of Information* (Cambridge: MIT Press, 1981) played perhaps the primary roles in leading epistemology to connect with science. Dretske, almost like Quine, allows the science to swamp the normative focus of epistemology. And the science that he uses, information theory, is, in my view, too far from the psychology of perception and perceptual belief to undergird epistemology. But his initiative, like Goldman's, was important.

perceptual beliefs also help provide a good epistemic route back to the subject matter. They do so when they yield reliably veridical perceptual states and perceptual beliefs. This route helps ground perceptual entitlement and commonly yields knowledgeable perceptual beliefs.

These displacements imply that neither the nature of representational mind nor the nature of empirical epistemic warrant can be explicated purely by reference to what goes on in the individual. Neither individuation of mental kinds nor empirical warrant is self-contained. Both constitutively depend on patterns of relations between individuals and the physical environment. These dependencies crucially engage unconscious, non-rational aspects of mind.

As indicated, broadly anti-individualistic views about mind were dominant in philosophy before the twentieth-century. The views about warrant were not, at least not since the early modern period. The explanation for this fact is probably more complex than anyone yet understands. I cite some factors.

The original and perhaps primary factor was confidence that Euclid's geometry modeled the structure of all knowledge or scientific-cognition. On the Euclidean model, there is a system of basic beliefs whose natures provide self-support. These beliefs support all other warranted beliefs. The starting points and inferential connections are supposed to be warranted by their natures. The model was often applied to empirical as well as mathematical and philosophical knowledge. In fact, in a persistent distortion, many prominent philosophers in the tradition took epistemically basic empirical beliefs to be approximately as secure from revision as beliefs in simple mathematical truths. Even since this delusion was exploded in mid-twentieth century, many philosophers remain in the grip of the Euclidean model. Most acknowledge that, under sufficient pressure from theoretical considerations, supported by other basic empirical beliefs, any given basic empirical belief can be rejected. But some cling to the idea that basic empirical beliefs are *prima facie* warranted by their natures. Others take perceptual contents to be intrinsic sources of *prima facie* warrant and to be analogs of Euclidean premises—except that they are not propositions and their contribution to warrant is *prima facie*.<sup>126</sup> Such positions prescribe doing epistemology by reflecting only on epistemic norms and individuals' mental states.<sup>127</sup> But reliability relations to the normal environment that are key elements in any empirical warrant are not coded into the natures of perceptual states, perceptual beliefs, or many inductive transitions.

<sup>126</sup> This analogy is thin. Perceptual states are indeed representational states that cause perceptual beliefs. They are the first representational states relevant to empirical knowledge. They are necessary for entitlement to perceptual beliefs. They bear normative relations to formation of perceptual beliefs. But nothing about the representational content of a perception *itself* provides any basis at all for the belief-worthiness of the belief. The representational content of a perceptual belief is nearly *identical* with part of the content of the perceptual state that underlies it, except for being in propositional form. The content's being *perceptual* can help explain belief-worthiness, but only if reliable relations to the environment are in place. These relations are not metaphysically determined by the very nature of the perception, as all warranting factors (including reliability) were supposed to be, on the Euclidean model.

<sup>127</sup> The peculiar view, naive realism, is a special case of this tendency. Like some of the more mainstream views that I have focused on, naive realism is out of step with what is known from the science of perception. For discussion, see 'Disjunctivism and Cognitive Psychology', *op. cit.* and 'Disjunctivism Again', *op. cit.*

A second factor is the human tendency to rely too much on familiar perspectives to set the limits of a subject matter. The tendency is strong in philosophy. It tends to hyper-intellectualize such subject matters as mind and knowledge, so that they are directly accessible to philosophical reflection. The focus on knowledge deluxe and insufficient reflection on perceptually based knowledge in young children have obscured how unsophisticated empirical knowledge can be.

A third factor is the idea that scepticism must be answered if epistemic warrant is to be obtained. The idea has dominated philosophical thinking since Descartes. Moore's on-point counter-initiative has been widely decried as obtusely irrelevant.

A fourth factor is a tendency to think of epistemic norms as always prescriptive—on an analogy with norms in ethics. Norms are standards for fulfilling functions, goals, or purposes. The most basic epistemic norms do not require any awareness of the norms or any ability to be guided by them. They apply not only to deliberative reasoning but to operations in partly unconscious psychological systems that are not under the control of the knower.

A fifth factor is failure to connect epistemic warrant constitutively to being conducive to knowledge and on a good route to truth. This failure arose only in mid-twentieth century.

A sixth factor is the practice of offering epistemic accounts of perceptual belief without having a scientifically grounded understanding of perception. This factor has dominated the history of philosophy. It has been less excusable over the last half-century, during which science has begun to understand perception in some depth. It is no longer defensible to do epistemology of perceptual belief, or philosophy of perception, without understanding perceptual psychology.

A scientific understanding of perception has made several contributions to the present account. I highlight two here.

One centers on the relation between veridicality and the nature of perceptual states. The science gives accounts of how both approximately accurate and illusory perceptual states are formed. It individuates perceptual states in terms of their veridicality conditions. It presupposes anti-individualism—which holds that what environmental features perceptual states attribute hinges partly on what environmental features figured prominently in the evolutionary pre-history of individuals with the perceptual systems. On the other hand, we know from biology that evolution does not directly select for accuracy. It selects for fitness to reproduce. Indeed, loss functions that help determine what perceptual states are generated from given stimulation are not set to maximize accuracy. They maximize avoidance of errors that lead to disadvantage or death. Although cases in which perceptual systems are approximately accurate are constitutive to determining the natures of the states, being *reliably* approximately accurate is not constitutive to the natures of the states. It is metaphysically possible that reliable radical inaccuracy, even in a content-determining environment, could yield environmental fitness. So reliable approximate veridicality is not constitutive of perceptual states or perceptual beliefs. Hence being epistemically warranted is not constitutive of perceptual beliefs.

A second contribution that a scientific understanding of perception makes is to elicit sources of our entitlements to perceptual beliefs. Anticipations and reliances in perceptual systems are calibrated to the probabilities of distributions of errors.

This fact shows how well perceptual systems respond to past errors and how reliable they in fact are, at least for a significant range of attributes in many normal circumstances. Actual perceptual states and perceptual beliefs seem almost as reliably veridical, for most types of attribution important to individuals' survival, as physical limitations allow. These points show in detail how perceptual systems enable perceptual beliefs to track approximate truth and gain approximate knowledge. Such systems monitor upshots of the same types of environment-perceptual-system connections that formed the attributives in the first place. The anticipatory dispositions and perceptual commitments coded into perceptual systems and perceptual-belief-forming systems are reliable in ways that contribute to entitlement. To be warranted, one need not be able to buttress perceptual beliefs with background thoughts. The source of warrant lies in pre-rational competencies that connect perceptual beliefs reliably to the world.

Mindless forces molded our perceptual and perceptual-belief-forming systems and made them epistemically good. Reason is not needed to justify the basic empirical products of representational systems. By fulfilling its very different functions, nature fulfilled some of reason's epistemic functions. Nature provided good routes to truth and knowledge for creatures with perceptual beliefs. Nature mindlessly formed an epistemic foundation for empirical reason.

## Bibliography

- Alston, W. (1976), 'Two Types of Foundationalism', *The Journal of Philosophy* 73, 165–185.
- Alston, W. (1983), 'What is Wrong with Immediate Justification?', *Synthese* 55, 73–95.
- Alston, W. (1989), *Epistemic Justification: Essays in the Theory of Knowledge* (Ithaca, New York: Cornell University Press).
- Alston, W. (2005), *Beyond "Justification": Dimensions of Epistemic Evaluation* (Ithaca, New York: Cornell University Press).
- Armstrong, D. (1973), *Belief, Truth, and Knowledge* (Cambridge: Cambridge University Press).
- Banks, M., J. Burge, and R. Held (2011), 'The Statistical Relationship between Depth, Visual Cues, and Human Perception', in M. Landy ed., *Sensory Cue Integration* (Oxford: Oxford University Press).
- Berker, S. (2013a), 'Epistemic Teleology and the Separateness of Propositions', *Philosophical Review* 122 (2013), 337–393.
- Berker, S. (2013b), 'The Rejection of Epistemic Consequentialism', *Philosophical Issues* 23, 363–387.
- Blakemore, C. (1970), 'The Range and Scope of Binocular Depth Discrimination in Man', *Journal of Physiology* 211, 599–622.
- Block, N. (2016), 'The Anna Karenina Principle and Skepticism about Unconscious Perception', *Philosophy and Phenomenological Research* 93, 452–459.
- BonJour, L. (1980), 'Externalist Theories of Empirical Knowledge', *Midwest Studies in Philosophy* 5, 53–73.
- BonJour, L. (1985), *The Structure of Empirical Knowledge* (Cambridge, Mass.: Harvard University Press).
- BonJour, L. (1998), *In Defense of Pure Reason* (London: Cambridge University Press).
- Burge, J. and W. Geisler (2014), 'Optimal Disparity Estimation in Natural Stereo Images', *Journal of Vision* 14, 1–18.

- Burge, J. and W. Geisler (2015), 'Optimal Speed Estimation in Natural Image Movies Predicts Human Performance', *Nature Communications* 6, Article number 7900.
- Burge, T. (1988), 'Individualism and Self-Knowledge', *The Journal of Philosophy* 85, 649–663, reprinted in T. Burge, *Cognition Through Understanding* (Oxford: Oxford University Press, 2013).
- Burge, T. (1997), 'Interlocution, Perception, and Memory', *Philosophical Studies* 86, 21–47.
- Burge, T. (2000), 'Frege on Apriority' in C. Peacocke and P. Boghossian eds., *New Essays on the A Priori*, (Oxford: Oxford University Press); reprinted in *Truth, Thought, Reason: Essays on Frege* (Oxford: Oxford University Press, 2005).
- Burge, T. (2003a), 'Perceptual Entitlement', *Philosophy and Phenomenological Research* 67, 503–548.
- Burge, T. (2003b), 'Reply to Stroud', in M. Hahn and B. Ramberg, *Reflections and Replies: Essays on the Philosophy of Tyler Burge* (Cambridge, Mass.: MIT Press).
- Burge, T. (2003c), 'Reply to Martin Davies', in M. Frapolli and E. Romero eds., *Meaning, Basic Self-Knowledge, and Mind: Essays on Tyler Burge* (Stanford: CSLI Publications).
- Burge, T. (2003d), 'Logic and Analyticity', *Grazer Philosophische Studien* 66, 199–249.
- Burge, T. (2005a), 'Disjunctivism and Perceptual Psychology', *Philosophical Topics* 33, 1–78.
- Burge, T. (2005b), *Truth, Thought, Reason: Essays on Gottlob Frege* (Oxford: Oxford University Press).
- Burge, T. (2007a), 'Postscript to "Belief De Re"', in *Foundations of Mind* (Oxford: Oxford University Press).
- Burge, T. (2007b), 'Predication and Truth', *The Journal of Philosophy*, 104, 580–608.
- Burge, T. (2007c), *Foundations of Mind: Philosophical Essays, Volume II* (Oxford: Oxford University Press).
- Burge, T. (2010), *Origins of Objectivity* (Oxford: Oxford University Press).
- Burge, T. (2011a), 'Epistemic Warrant: Humans and Computers', in *Cognition Through Understanding: Philosophical Essays, Volume III* (Oxford: Oxford University Press, 2013).
- Burge, T. (2011b), 'Origins of Perception', *Disputatio* 4, 1–38.
- Burge, T. (2011c), 'Steps Toward Origins of Propositional Thought', *Disputatio* 4, 39–67.
- Burge, T. (2011d), 'Disjunctivism Again', *Philosophical Explorations* 14, 43–80.
- Burge, T. (2013), *Cognition through Understanding: Philosophical Essays, Volume III* (Oxford: Oxford University Press).
- Burge, T. (2014a), 'Reply to Rescorla and Peacocke: Perceptual Content in Light of Perceptual Constancies and Biological Constraints', *Philosophy and Phenomenological Research* 88, 485–501.
- Burge, T. (2014b), 'Perception: Where Mind Begins', *Philosophy* 89, 385–403.
- Burge, T. (2014c), 'Reply to Block: Adaptation and the Upper Border of Perception', *Philosophy and Phenomenological Research* 89, 573–583.
- Camp, E. (2007), 'Thinking with Maps', *Philosophical Perspectives* 21, 145–182.
- Campbell, J. (2002), *Reference and Consciousness* (Oxford: Clarendon Press).
- Casullo, A. (2007), 'What Is Entitlement?', *Acta Analytica* 22, 267–279.
- Chisholm, R.M. (1966), *Theory of Knowledge*, 1<sup>st</sup> edition (Englewood Cliffs, N.J.: Prentice Hall).
- Chisholm, R.M. (1977), *Theory of Knowledge*, 2nd edition (Englewood Cliffs, New Jersey: Prentice Hall).
- Christensen, M., L. Kristiansen, J. Rowe, and J. Nielsen (2008), 'Action-blindsight in Healthy Subjects after Transcranial Magnetic Stimulation', *PNAS* 105, 1353–1357.
- Clarke, T. (1972), 'The Legacy of Skepticism', *The Journal of Philosophy* 69, 754–769.
- Cohen, S. (1984), 'Justification and Truth', *Philosophical Studies* 46, 279–295.
- Cohen, S. (2000), 'Contextualism and Skepticism', *Philosophical Issues* 10, 94–107.

- Cohen, S. (2002), 'Basic Knowledge and the Problem of Easy Knowledge', *Philosophy and Phenomenological Research* 65, 309–329.
- Cohen, S. (2005), 'Why Basic Knowledge is Easy Knowledge', *Philosophy and Phenomenological Research* 70, 417–430.
- Cohen, S. (2010), 'Bootstrapping, Defeasible Reasoning, and A Priori Justification', *Philosophical Perspectives* 24, 141–159.
- Conee, E. and R. Feldman (1998), 'The Generality Problem for Reliabilism', *Philosophical Studies* 89, 1–29.
- Cruz J., and J. Pollock (2004), 'The Chimerical Appeal of Epistemic Externalism', in R. Schantz ed., *The Externalist Challenge* (Berlin: De Gruyter).
- Danckert, J. and Y. Rossetti (2005), 'Blindsight in Action: What Can the Different Sub-Types of Blindsight Tell Us about the Control of Visually Guided Actions?', *Neuroscience & Biobehavioral Reviews* 29, 1035–1046.
- Danto, A. (1968), *Analytical Philosophy of Knowledge* (Cambridge: Cambridge University Press).
- D'Antona, D., J. Perry, and W. Geisler (2013), 'Humans Make Efficient Use of Natural Image Statistics', *Journal of Vision* 13, 1–13.
- Davidson, D. (1983), 'A Coherence Theory of Truth and Knowledge' in *Subjective, Intersubjective, Objective* (Oxford: Clarendon Press, 2001).
- Davies, M. (2000), 'Externalism and Armchair Knowledge', in P. Boghossian and C. Peacocke eds., *New Essays on the A Priori* (Oxford: Oxford University Press).
- DeRose, K. (1995), 'Solving the Skeptical Problem', *The Philosophical Review* 104, 1–52.
- Dickie, I. (2011), 'Visual Attention Fixes Demonstrative Reference by Eliminating Referential Luck' in C. Mole, D. Smithies, and W. Wu eds., *Attention: Psychological and Philosophical Essays* (Oxford: Oxford University Press).
- Dretske, F. (1981), *Knowledge and the Flow of Information* (Cambridge: MIT Press).
- Ernst, M. and M. Banks (2002), 'Humans Integrate Visual and Haptic Information in a Statistically Optimal Fashion', *Nature* 415, 429–433.
- Foley, R. (1985), 'What's Wrong with Reliabilism?', *The Monist* 68, 188–202.
- Frisby, J. and J. Stone (2010), *Seeing* (Cambridge, Mass.: MIT Press).
- Fumerton, R. (1995), *Meta-epistemology and Skepticism* (Totowa, New Jersey: Rowman and Littlefield).
- Geisler, W. (2003), 'Ideal Observer Analysis', in L. Chalupa and J. Werner eds., *The Visual Neurosciences* (Boston: MIT Press).
- Geisler, W. (2008), 'Visual Perception and the Statistical Properties of Natural Scenes', *Annual Review of Psychology* 59, 167–192.
- Geisler, W. and R. Diehl (2003), 'A Bayesian Approach to the Evolution of Perceptual and Cognitive Systems', *Cognitive Science* 27, 379–402.
- Geisler, W., J. Perry, B. Super, and D. Gallogly (2001), 'Edge Cooccurrence in Natural Images Predicts Contour Grouping Performance', *Vision Research* 41, 711–724.
- Ghosi, A. (2015), 'Dimensionality Reduction: A Short Tutorial', [http://www.math.uwaterloo.ca/~aghodsib/courses/f06stat890/readings/tutorial\\_stat890.pdf](http://www.math.uwaterloo.ca/~aghodsib/courses/f06stat890/readings/tutorial_stat890.pdf) (sourced August 2015).
- Ginet, C. (1975), *Knowledge, Perception, and Memory* (Dordrecht: D. Reidel).
- Goldman, A. (1979), 'What is Justified Belief?', in G. Pappas ed., *Justification and Knowledge* (Dordrecht: Reidel).
- Goldman, A. (1986), *Epistemology and Cognition* (Cambridge, Mass.: Harvard University Press).
- Goldman, A. (1988), 'Strong and Weak Justification', *Philosophical Perspectives* 2, 51–69.

- Goldman, A. (2001), 'The Unity of the Epistemic Virtues', in A. Fairweather and L. Zagzebski eds., *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility* (Oxford: Oxford University Press).
- Goldman, A. (2008), 'Immediate Justification and Process Reliabilism', in Q. Smith ed., *Epistemology: New Essays* (Oxford: Oxford University Press).
- Goldman, A. 'Reliabilism', *Stanford Encyclopedia of Philosophy*, online.
- Greco, J. (2006), 'Virtue, Luck and the Pyrrhonian Problematic', *Philosophical Studies*, 130, 9–34.
- Hawthorne, J. (2004), *Knowledge and Lotteries* (Oxford: Oxford University Press).
- Huemer, M. (2001), *Skepticism and the Veil of Perception* (Lanham, Md.: Rowman & Littlefield).
- Huemer, M. (2006), 'Phenomenal Conservatism and the Internalist Intuition', *American Philosophical Quarterly* 43, 147–158.
- Kung, P. (2010), 'On Having No Reason: Dogmatism and Bayesian Confirmation', *Synthese* 177, 1–17.
- Lewis, D. (1996), 'Elusive Knowledge', *Australasian Journal of Philosophy* 74, 549–567.
- Luthra, Y. (2013), Embodied Rational Agency, UCLA Dissertation, (2013), chapter 6—retrieved from <https://escholarship.org/uc/item/4q92g4w6>.
- Majors, B. (2015), 'What Entitlement Is,' *Acta Analytica*, 30, 363–87.
- Maloney, L. and P. Mamassian (2009), 'Bayesian Decision Theory as a Model of Human Perception: Testing Bayesian Transfer', *Visual Neuroscience* 26, 147–155.
- Moon, A. (2012), 'Three Forms of Internalism and the New Evil Demon Problem', *Episteme* 9, 345–360.
- Moore, G.E. (1939), 'Proof of an External World', *Proceedings of the British Academy* 25, 273–300.
- Moore, G.E. (1953), *Some Main Problems of Philosophy* (London: George Allen & Unwin Ltd).
- Moore, G.E. (1993), 'Moore's Paradox', in T. Baldwin ed., *G. E. Moore: Selected Writings* (New York: Routledge).
- Moretti, L. (2015), 'In Defense of Dogmatism', *Philosophical Studies* 172, 261–282.
- Nozick, R. (1981), *Philosophical Explanations* (Cambridge, Mass.: Harvard University Press).
- Peacocke, C. (2004), *The Realm of Reason* (Oxford: Oxford University Press).
- Peacocke, C. (2015), 'Magnitudes: Metaphysics, Explanation, and Perception', in D. Moyal-Sharrock, V. Munz, and A. Coliva eds., *Mind, Language and Action: Proceedings of the 36th International Wittgenstein Symposium* (Berlin: de Gruyter).
- Perenin, M. and Y. Rossetti (1996), 'Grasping Without Form Discrimination in a Hemianopic Field', *Neuroreport* 7, 793–797.
- Philips, I. (2018), 'Unconscious Perception Reconsidered', *Analytic Philosophy* 59, 471–514.
- Philips, I. and N. Block (2016), "Unconscious Seeing"—A Debate', in B. Nanay ed., *Current Controversies in Philosophy of Perception* (London: Routledge).
- Plantinga, A. (1993), *Warrant: The Current Debate* (Oxford: Oxford University Press).
- Pollock, J. (1974), *Knowledge and Justification* (Princeton: Princeton University Press).
- Pollock, J. (1984), 'Reliability and Justified Belief', *Canadian Journal of Philosophy* 14, 103–114.
- Pollock, J. and J. Cruz (1999), *Contemporary Theories of Knowledge* (Lanham, Maryland: Rowman and Littlefield).
- Pryor, J. (2000), 'The Skeptic and the Dogmatist', *Nous* 34, 517–549.
- Pryor, J. (2001), 'Highlights of Recent Epistemology', *British Journal of the Philosophy of Science* 52, 95–124.
- Pryor, J. (2004), 'What is Wrong with Moore's Argument?' *Philosophical Issues* 14, 349–378.



- Pryor, J. (2005), 'There is Immediate Justification' in M. Steup and E. Sosa eds., *Contemporary Debates in Epistemology* (Oxford: Blackwell).
- Quine, W.V. (1960), *Word and Object* (Cambridge, Mass.: MIT Press).
- Quine, W.V. (1969), 'Epistemology Naturalized', in *Ontological Relativity* (New York: Columbia University Press).
- Quine, W.V. and J. Ullian (1970), *The Web of Belief* (New York: Random House).
- Quinton, A. (1973), *The Nature of Things* (London: Routledge & Kegan Paul).
- Reid, T. (2002), *Essays on the Intellectual Powers of Man* (1785) (Edinburgh.: Edinburgh University Press).
- Rescorla, M. (2015), 'Bayesian Perceptual Psychology', in M. Matthen ed., *The Oxford Handbook of the Philosophy of Perception* (Oxford: Oxford University Press).
- Rescorla, M. (2016), 'Bayesian Sensorimotor Psychology', *Mind and Language* 31, 3–36.
- Scanlon, T. (1998), *What We Owe Each Other* (Cambridge, Mass.: Harvard University Press).
- Scanlon, T. (2014), *Being Realistic About Reasons* (Oxford: Oxford University Press).
- Sellars, W. (1963), 'Empiricism and the Philosophy of Mind' in *Science, Perception, and Reality* (London, Routledge & Kegan Paul).
- Shadmehr, R. and S. Mussa-Ivaldi (2012), *Biological Learning and Control* (Cambridge, Mass.: MIT Press).
- Silins, N. (2008), 'Basic Justification and the Moorean Response to the Skeptic', in T. Gendler and J. Hawthorne eds., *Oxford Studies in Epistemology* 2, 108–142.
- Silins, N. (2012), 'Explaining Perceptual Entitlement', *Erkenntnis* 76, 243–261.
- Smithies, D. (2011), 'What is the Role of Consciousness in Demonstrative Thought?', *The Journal of Philosophy* 108, 5–34.
- Sosa, E. (1994), 'Philosophical Skepticism and Externalist Epistemology', *Proceedings of the Aristotelian Society*, Supplementary Volume 68, 263–290.
- Sosa, E. (1997), 'How to Resolve the Pyrrhonian Problematic: A Lesson from Descartes', *Philosophical Studies* 85, 229–249.
- Sosa, E. (2007), *A Virtue Epistemology* (Oxford: Clarendon Press).
- Sosa, E. (2009), *Reflective Knowledge* (Oxford: Clarendon Press).
- Sosa, E. (2011), *Knowing Full Well* (Princeton: Princeton University Press).
- Sperandio, I., and P. Chouinard (2015), 'The Mechanisms of Size Constancy' *Multisensory Research* 28, 253–283.
- Steup, M. (2004), 'Internalist Reliabilism', *Philosophical Issues* 14, 403–425.
- Stroud, B. (1984), *The Significance of Philosophical Scepticism* (Oxford: Clarendon Press).
- Stroud, B. (1989), 'Understanding Human Knowledge in General' in M. Clay and K. Lehrer eds., *Knowledge and Skepticism* (Boulder: Westview).
- Stroud, B. (1994), 'Scepticism, "Externalism", and the Goal of Epistemology', *Proceedings of the Aristotelian Society*, Supplementary Volumes 68, 291–307.
- Vogel, J. (1990), 'Cartesian Skepticism and Inference to the Best Explanation', *Journal of Philosophy* 87, 658–666.
- Vogel, J. (2000), 'Reliabilism Leveled', *The Journal of Philosophy* 97, 602–623.
- Weatherson, B. (2007), 'The Bayesian and the Dogmatist', *Proceedings of the Aristotelian Society* 107, 169–185.
- Weatherson, B. (2014), 'Easy Knowledge and Other Epistemic Virtues', online <http://brian.weatherson.org/ekoev.pdf%20sourced%202014>.
- Wedgewood, R. (2002), 'Internalism Explained', *Philosophy and Phenomenological Research* 65, 349–369.
- Wedgewood, R. (2011), 'Primitively Rational Belief-Forming Processes', in A. Reisner and A. Steglich-Petersen eds., *Reasons for Belief* (Cambridge: Cambridge University Press).

- White, R. (2006), 'Problems for Dogmatism', *Philosophical Studies* 131, 525–557.
- Whitwell, R., C. Streimer, D. Nicole, and M. Goodale (2011), 'Grasping the Non-conscious: Preserved Grip Scaling to Unseen Objects for Immediate but not Delayed Grasping Following a Unilateral Lesion to Primary Visual Cortex', *Vision Research* 51, 908–924.
- Whitwell, R., A. Milner, C. Cavina-Pratesi, M. Barat, and M. Goodale (2015a), 'Patient DF's Visual Brain in Action', *Vision Research* 110, 265–276.
- Whitwell, R., I. Sperandio, G. Buckingham, P. Chouinard, and M. Goodale (2015b), 'Evidence for a Functional and Anatomical Dissociation in the Use of Size Constancy for Perceptual Report and Goal-directed Grasping' (abstract), *Journal of Vision* 15, 187.
- Williamson, T. (2001), *Knowledge and Its Limits* (Oxford: Oxford University Press).
- Wolpert, D. and M. Landy (2012), 'Motor Control is Decision-Making', *Current Opinion in Neurobiology* 22, 1–8.
- Wright, C. (2002), '(Anti)-Sceptics Simple and Subtle: Moore and McDowell', *Philosophy and Phenomenological Research* 65, 330–48.
- Wright, C. (2004), 'Warrant for Nothing (and Foundations for Free)?', *Aristotelian Society*, Supplementary Volume 78, 167–2012.
- Wright, C. (2007), 'The Perils of Dogmatism', in S. Nuccetelli and G. Seay eds., *Themes from G. E. Moore: New Essays in Epistemology and Ethics* (Oxford: Oxford University Press).
- Yuille, A. and H. Bulthoff (1996), 'Bayesian Decision Theory and Psychophysics' in D. Knill and W. Richards eds., *Perception as Bayesian Inference* (Cambridge: Cambridge University Press).
- Zardini, E. (2014), 'Confirming the Less Likely, Discovering the Unknown', in D. Dodd and E. Zardini eds., *Scepticism and Perceptual Justification* (Oxford, Oxford University Press).