

Perceptual Justification and the Phenomenology of Experience

Jorg DhiptaWillhofs

UCL

Submitted for the Degree of PhD

I, Jorg Dhipta Willhoft, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

This thesis seeks to provide an explanation of what I call the Basic Principle about Perceptual Justification which states that if a subject *S* has a perceptual experience as of a mind-independent object *x* being *F* (or in which it appears to him as if an *x* is *F*), and forms the belief that an *x* is *F* on the basis of having an experience of this phenomenological sort, then (perhaps provided certain further conditions obtain) *S*'s belief that an *x* is *F* is *prima facie* justified for *S*. I distinguish between two conceptions of epistemic justification. Roughly, on an objective conception, a subject *S* has a justified belief that *p* if he bases this belief on grounds that entail or make likely the truth of *p*, while on a subjective conception a subject *S* has a justified belief that *p* if he forms this belief on the basis of his occupying a perspective from which a situation obtains that entails or makes likely the truth of *p*. I argue that the truth of the Basic Principle can be derived, in part, from facts about the phenomenal character of perceptual experience. In particular, I argue that the Basic Principle can be explained by saying that the subject's perceptual experience can provide him with justification for believing that an *x* is *F* in the subjective sense and that it does so, in part, in virtue of its phenomenal character. I also address the question of whether perceptual experiences can provide us with immediate justification for believing propositions about our environment, that is, with justification that does not depend on our having independent justification for believing other propositions such as the proposition that perceptual experiences are generally reliable. To this end, I consider the so-called problem of easy knowledge and argue that the issues concerning this problem should not compel us into thinking that perceptual justification cannot be immediate.

Contents

Acknowledgements	5
Introduction	7
Phenomenology and the Basic Principle	7
Bonjour's Challenge.....	15
Motivating a Phenomenological Approach.....	26
Outline of the Thesis	29
Chapter 1	32
Introduction	32
Global Reliabilism	37
Bergman's Proper Function Account.....	45
Burge's Anti-Individualist Account.....	51
Evaluating Bergman's and Burge's Views	60
Chapter 2	67
Introduction	67
Foley's Conception of an Epistemically Rational Belief.....	72
McDowell's Conception of Perceptual Justification.....	82
E-Representing the Obtaining of a Mind-Independent Fact	100
Chapter 3	107
Introduction	107
McDowell's Response to Scepticism about the External World	112
Wright's Response to McDowell's Anti-Scepticism	116
Perceptual Knowledge and Conclusive Warrants	127
Chapter 4	144
Introduction	144
The Closure Problem.....	150
The Bootstrapping Problem	161
Knowing vs. Claiming to Know.....	174
Bibliography.....	180

Acknowledgements

I would like to thank all my supervisors over the years, Mark Kalderon, Christopher Peacocke, Paul Snowdon, and Jose Zalabardo, for helpful comments on different parts of this thesis and for guiding me through various issues in epistemology and the philosophy of perception. Any mistakes or misconceptions that the thesis may contain are, of course, entirely my responsibility. I would also like to thank the Philosophy Department at UCL for their financial support and Rory Madden and Jose for helping me to secure a much-needed extension to my original submission deadline.

Finishing this thesis would not have been possible if I had not received the help of friends and family in taking care of my two sons, Han and Tang. So a warm and heartfelt thank you goes to Midori Ainoura, Sally Holyland, Takatomo Kashiwabara, Katie Kelly, Michelle King, Stephen Ooi, Anne Stranne-Petersen, Oliver Willhoft, Udo Willhoft, and Ana Yerill. I cannot thank enough my mother, Tuti Willhoft, and my mother-in-law, Ellen King, both of whom stayed with us for extended periods of time, took care of us, showered the kids with love, and kept things together for us.

I owe the most to my wife, Laura King. She has made great sacrifices and worked tirelessly to allow me to pursue my goals. In return, I have stretched her patience to the limit by missing several personal deadlines. But our ‘shared custody’ arrangement, in which I have the kids during the week and she has them during the weekend, is now at an end and we can finally be a proper family. Without Laura, little of what I do would be possible and none of it would be worthwhile. And so, with gratitude and apologies but above all with love, I dedicate this thesis to her.

To Laura

Introduction

Do perceptual experiences contribute to the epistemic justification of beliefs about the external world? And if they do, how and under what circumstances do they do so? These are the questions that I will be addressing in this thesis. I will begin with some natural and intuitive thoughts, first, about the nature of perceptual experiences and, second, about their epistemic role in regards to external world beliefs.

Phenomenology and the Basic Principle

We can begin with an insight by P.F. Strawson, who claims that “mature experience (in general) presents itself as, in the Kantian phrase, an immediate consciousness of the existence of things outside us.” (Strawson, 1979, p. 47) Strawson’s claim here is about how perceptual experience presents itself to us when we reflect on, or introspect, its nature, and it is a claim that can be divided into two parts. The first is that, on reflection, a perceptual experience presents itself as a conscious awareness of certain ‘things’ or objects. The second is that these objects are ‘things outside us’ or mind-independent in the sense that they exist independently of our conscious awareness of them. If we take the first part first, what exactly is involved in the (seeming) conscious awareness of objects that we attribute to perceptual experience? A helpful formulation of what is involved here is provided by John Searle in the following passage:

If, for example, I see a yellow station wagon in front of me, the experience I have is directly of the object. It doesn't just "represent" the object, it provides direct access to it. The experience has a kind of directness, immediacy and involuntariness which is not shared by a belief I might have about the object in its absence. (Searle, 1983, pp. 45-6)

Searle's aim in this passage is to capture certain features of perceptual experiences that he designates with the use of words like 'directness' and 'immediacy'. Thus the seeming conscious awareness that we find in perceptual experience when we reflect on them is a kind of direct presentation to us of these objects; in having a perceptual experience, it is, to put it differently, as if we are directly confronted with certain objects as *being right there* before us (an idea that is also captured by Strawson's formulation of experience as involving a seeming consciousness of the *existence* of things outside us) and, though this is not mentioned by Searle in this passage, as *having certain properties*. Call this feature the *presentational character* of experience. In having this feature, perceptual experiences essentially differ from other mental states such as thoughts or beliefs. For example, in thinking about a yellow station wagon, it need not be the case that it is to one as if a yellow station wagon is right there in front of one.

Consider now the second part of Strawson's claim. The idea here was that perceptual experience, on introspection, presents itself as a conscious awareness of *mind-independent* objects. In making this claim, part of Strawson's point is to say that a description of perceptual experience which is 'faithful' to 'experience as we actually enjoy it' (Strawson, 1979, p. 43) will essentially involve the use of concepts for mind-independent objects. Thus when asked to give a description of his current

visual experience, a person might say, for example, that he sees ‘the red light of the setting sun filtering through the black and thickly clustered branches of the elms...the dappled deer gazing in groups on the vivid green grass...’ (ibid, p. 43) Of course, this description of the visual experience in terms of the person’s seeing certain mind-independent objects involves a commitment to the existence of those objects. Strawson insists, however, that even if one were to shed this kind of commitment, a faithful description of the person’s visual experience would still involve the very same concepts for mind-independent objects involved in the original description. Thus the person might say instead that he ‘had a visual experience such as it would have been natural to describe by saying that [he] saw [the red light of the setting sun...].’ (ibid, pp.43-4)

Call this feature the *world-directedness* of perceptual experience. Together the presentational character of a perceptual experience and its world-directedness yield a description of perceptual experience as a seeming confrontation with the external world – where how things are in the world seems to be made manifest to one in having the visual experience. An important point to note is that we find both of these two features in perceptual experiences that amount to genuine perceptions of an object in the external world as well as in experiences that do not amount to genuine perceptions. For example, when we see a red ball, reflection on this experience would present this experience as an immediate conscious awareness of a red ball that exists independently of our conscious awareness of it. However, these observations apply equally to non-veridical experiences such as hallucinations. Thus we can imagine being in a state that, from the inside, would appear just to be the same as a genuine veridical perception of a red ball but where there is in fact no red ball before us to be seen. Still, even if there is no relevant object to be seen, the

experience – when described from a subjective point of view – presents itself as an immediate conscious awareness of a red ball.

These very natural thoughts are thoughts essentially about the *phenomenal character* of perceptual experiences. They are meant to capture *what it is like* for us, from our own subjective point of view, to undergo such experiences. Thus, phenomenologically speaking, veridical (visual) perceptions as well as (visual) hallucinations seem to involve a conscious awareness of mind-independent objects and some of their properties.

Apart from these natural thoughts about the phenomenal character of perceptual experience, there are also some very natural thoughts about the epistemic role of perceptual experiences. Thus it is very natural to think that many of our beliefs about our surroundings that we form on the basis of our perceptual experiences are epistemically justified for us. Suppose, for example, that as a result of seeing a red ball one forms the belief that there is a red ball in front one. Under normal circumstances, a belief of this kind seems to be entirely justified. To be sure, one's justification here may be defeated or undermined at a later point. For example, one may acquire evidence that epistemically supports the belief that one is not genuinely perceiving, in which case the justification provided by the perceptual experience may be undermined. But in the absence of any countervailing evidence, it seems to be epistemically appropriate for us to hold our perceptually based beliefs. Perceptions of mind-independent objects, therefore, seem to provide us with *prima facie* justification for believing certain propositions about those objects.

These thoughts seem to apply equally to perceptual experiences that do not amount to fully veridical perceptions of mind-independent objects. Intuitively, these too seem to have the capacity to provide us with some *prima facie* justification for

beliefs about our surroundings. If, for example, one unwittingly suffers from a highly deceptive hallucinatory experience in which it appears to one as if there is a red ball in front of one when there is in fact no red ball for one to be seen and on the basis of having this experience forms the belief that there is such an object in front of one, then this belief too may be entirely justified for one. Of course, such a belief may well be false. Still, in the absence of reasons for thinking that one is in fact suffering from a hallucination, it seems to be entirely appropriate, from an epistemic point of view, for one to believe that there is a red ball in front of one.

These intuitive thoughts about the phenomenal character and epistemic role of perceptual experiences can be combined to formulate the following principle about perceptual justification:

If a subject S has a perceptual experience as of a mind-independent object x being F (or in which it appears to him as if an x is F), and forms the belief that an x is F on the basis of having an experience of this phenomenological sort, then (perhaps provided certain further conditions obtain) S's belief that an x is F is prima facie justified for S.

I will refer to this principle as the *Basic Principle about Perceptual Justification* (or the Basic Principle for short). The antecedent of this principle is to be taken as making a claim about a subject enjoying a perceptual experience with the kind of phenomenal character that was described a moment ago. Thus a perceptual experience as of a mind-independent object x being F (or in which it appears to one as if an x is F) is an experience which involves a seeming conscious awareness of a mind-independent object x instantiating the property F.

The first point to note about this principle is that, stated as it is, it records a certain intuition about when a subject can have a prima facie justified belief about his surroundings. The intuition is that when a subject enjoys a perceptual experience of an x being F and forms the belief that an x is F on the basis of it then that belief is prima facie justified for him. The Basic Principle as such, however, does not have any implications on what makes such a belief epistemically justified for a subject; nor does it have any implications on what the role is of the perceptual experience in accounting for the justification of the belief. It also leaves open the possibility that further specific conditions have to obtain, other than the presence of a relevant sort of perceptual experience, in order for a belief based on this experience to have the status of being prima facie justified for one.

Another point to note is that this principle is a principle about what is usually referred to as *doxastic justification*. It thus states certain conditions under which a particular belief of a subject can have the property of being justified for him. In general we can say that for a subject's belief that p to have this property is for the subject to possess adequate grounds in support of believing that p and for him to form the belief that p on the basis of these grounds. The notion of doxastic justification so understood is different from the notion of *propositional justification*. The latter notion picks out a property that a proposition rather than a belief can have for a subject. Thus one can have propositional justification for believing a proposition p if one possesses adequate grounds in support of believing that p without thereby having to have a justified belief that p. This could be, for example, when one does not in fact form the belief that p despite having adequate grounds for believing that p, or when one does have the belief that p without however having based the belief on the grounds one possesses.

We can of course formulate the Basic Principle in terms of the notion of propositional justification. The result would be the following:

If a subject S has a perceptual experience as of a mind-independent object x being F (or in which it appears to him as if an x is F), then (perhaps provided certain further conditions obtain) S has some prima facie justification for believing that an x is F.

This thesis will consider various explanations of why the Basic Principle in either of these two forms is true. However, of particular interest will be the relation between this principle and the phenomenal character of perceptual experience. One question to be addressed is thus whether the truth of the Basic Principle can be derived from facts about what it is like for a subject to undergo the perceptual experience in such way that we can say that the perceptual experience contributes to the justification of his belief in virtue of its having the phenomenal character that it does. I will refer to accounts that attempt to explain the Basic Principle, at least in part, by appeal to facts about the phenomenal character of experiences as *phenomenological accounts of, or phenomenological approaches to, the Basic Principle*. An endorsement of such an account has recently been given by James Pryor. Thus Pryor states:

My view is that our perceptual experiences have the epistemic powers...they have because of what the phenomenology of perception is like.

I think there's a distinctive phenomenology: the feeling of *seeming to ascertain* that a given proposition is true. This is present when the way a mental episode represents its content makes it feel as though, *by enjoying that*

episode, you can *thereby just tell* that that content obtains...When you have a perceptual experience of your hands, that experience makes it feel as though *you can just see* that hands are present. It feels as though hands are being shown or revealed to you. (Pryor, 2004, pp. 356-7)

Pryor makes a number of claims in this passage that need not be adopted by all proponents of a phenomenological approach. For one, he assumes that perceptual experiences are mental states with representational contents and that it is in virtue of their having such contents (and the way these contents are represented in experience) that perceptual experiences have the phenomenal character that they do. Moreover, he seems to think that the phenomenal character of perceptual experiences involves different kinds of 'feelings' such as the feeling of seeming to ascertain a proposition or the feeling that one can just tell that a proposition is true. Finally, he believes that it is in virtue of these phenomenological facts that perceptual experiences have the capacity to provide us with justification for believing certain propositions about the external world.

None of these claims, however, are essential to the phenomenological approach. Thus the kind of *intentionalism* about the nature of perceptual experience that Pryor assumes here is only one of several competing accounts of what constitutes the phenomenal character of experience. Furthermore, we might have doubts about the contention that the phenomenal character of a perceptual experience involves the kinds of feeling that Pryor describes and, more generally, the contention that it involves any feelings at all. At the end of the passage, though, Pryor suggests that it is because in having the relevant perceptual experience it feels to one as if hands are being shown or revealed that such an experience can provide one with justification

for believing that there are hands in front of one. If we shed the notion that there is a kind of ‘feeling’ involved in having the perceptual experience, we can simply take this to reflect the features of the phenomenal character of experience we identified earlier, namely its presentational character and world-directedness. From the subjective point of view, when one sees one’s hands, it is to one as if there are hands right before one. The suggestion is then that it is in virtue of the perceptual experience having these phenomenal features that it provides one with justification for believing that there are hands in front of one.

The question of course is why we should think that this is the case. And more fundamentally: what exactly is it to say that it is *in virtue of* the phenomenal properties of a perceptual experience that this experience can provide one with justification for believing propositions about the external world? Pryor himself does not provide us with any answers to these questions, at least not in the foregoing passage. In what follows, I want to approach these questions by considering a particular challenge that we might pose to the proponent of a phenomenological approach.

Bonjour’s Challenge

This challenge is expressed by Laurence Bonjour in the following passage:

It is easy to see how the way in which material objects seem to be simply presented or, as one might even be tempted to say, given in perceptual experience could lead to the view that is usually (and plausibly) ascribed to at least the most naive level of common sense, namely that there is no problem

at all about the justification or indeed truth of the resulting beliefs. But...this presentational character of experience has to do with the way in which physical objects are represented or depicted in experience, but has no obvious bearing on whether such representations or the beliefs that reflect them are true. A presentational representation is no doubt more vivid, more striking, in something like the way in which a picture is more compelling than a merely verbal description. But pictures are just as capable of being mistaken as anything else, and so the pictorial character of a representation seems to be simply irrelevant to the issue of justification; my suggestion is that we have so far seen no clear reason not to say the same thing about the intuitively presentational character of perceptual experience. (Bonjour, *In Search of Direct Realism*, 2004, pp. 354-5)

The first point to note about this passage is that the analogy to pictures that Bonjour draws here is entirely misguided. As we have seen, a phenomenological approach to explaining the Basic Principle claims that it is, at least in part, in virtue of the phenomenal properties of a perceptual experience that the experience provides us with justification for believing a proposition about the external world. These phenomenal properties are properties of what it is like for a subject to undergo the perceptual experience. They belong to the general class of properties of what it is like for a subject to be in a conscious state. Such properties, however, are simply not attributable to pictures; that is, there is nothing it is like for a subject to be a picture. (Of course, there is the property of what it is like for a subject to be consciously aware of a picture but this is irrelevant to the issue at hand) So whatever kind of 'presentational character' a picture can have it is not the same kind as the

presentational character that we have identified as a feature of the phenomenal character of experience. Hence, even if there are reasons for doubting that the presentational character of a picture has any relevance at all with respect to the question of whether such pictures provide us with any justification for believing propositions about the world, we should not expect that these problems apply to perceptual experiences and their phenomenal character.

However despite the fact that the analogy to pictures is misguided, Bonjour still raises an important point in this passage. His central claim is that the presentational character of a perceptual experience, and presumably its phenomenal character as a whole, has no obvious bearing on the truth of the propositions that this experience purportedly provides us with justification for believing. And if it has no obvious bearing on the truth of the relevant propositions, then there is no (obvious) reason for thinking that it is in virtue of its phenomenal character that a perceptual experience can provide us with justification for believing those propositions to be true. Of course, even if there are no obvious reasons for thinking this, this does not mean that there are no such reasons or that such reasons cannot be discovered upon deeper reflection.

How then could a proponent of the phenomenological approach respond to Bonjour's challenge? Before we can answer this question, we first need to be clear about how exactly this challenge should be understood. The questions we need to address at this point are these: in what way should the phenomenal character of a perceptual experience 'bear on the truth' of the propositions for which such an experience provides us with justification for believing if it is in virtue of its phenomenal character that the experience has the given epistemic properties? And why should the phenomenal character of a perceptual experience bear on the truth of

the relevant propositions in this way in order for the experience to provide us with justification for believing those propositions? Finally, why should we think that the phenomenal character of perceptual experiences does not bear on the truth of the relevant propositions in the required way?

These questions need to be answered against the backdrop of certain assumptions about the concept of epistemic justification in general. One basic assumption about this concept is that it is essentially connected in some way with the concept of truth. This assumption in turn can be seen as being grounded in the general assumption that a belief essentially aims at truth. In fulfilling this aim, one will believe a proposition *p* only if *p* is true. Given this aim of belief, we can then understand the concept of epistemic justification as a specific kind of evaluative notion: a given belief is epistemically justified only if it is in some relevant sense appropriate with respect to the aim of believing a proposition *p* only if *p* is true. A theory of epistemic justification will therefore have to attempt to spell out the conditions under which it is appropriate in the relevant sense to hold a given belief. Following Stewart Cohen¹, we can assume that there are two general approaches we can take towards spelling out these conditions, one of which yields an objective notion of epistemic justification, while the other yields a subjective notion.

Take the former notion first. Suppose a belief that *p* is based on a ground *g*, such as a perceptual experience. We can then say that *g* provides one with justification for believing that *p* only if forming the belief that *p* on the basis of *g* reliably results in that belief's being true. Another way of formulating an objective notion of epistemic justification is in terms of belief-forming processes rather than grounds for believing. Thus, a belief-forming process produces justified beliefs only if this process reliably

¹ (Cohen, *Justification and Truth*, 1984)

produces beliefs that are true. Of course, the objective notion of justification understood in this way raises the question as to the degree of reliability that is required for there to be epistemic justification for believing that p, but we can set this question aside here. At the extreme end of objective conceptions of epistemic justification, a ground g provides one with justification for believing that p only if forming the belief that p on the basis of g necessarily results in the belief's being true. (Or, if we prefer a formulation in terms of processes: a belief-forming process pr produces justified beliefs only if pr necessarily produces beliefs that are true)

Alternatively, on a subjective notion of justification, we can think of epistemic justification not as a matter of one's having a ground g which is in fact reliably connected with the truth of p but rather as a matter of how things are from the subject's own perspective. Here, we can initially think of a perspective as being constituted by a subject's beliefs. Thus Cohen himself suggests as a version of the subjective notion of epistemic justification that a subject has justification for believing that p only if he has an 'impeccable belief' that a proposition q distinct from p obtains and that q makes likely the truth of p. (Cohen, *Justification and Truth*, 1984, p. 285) In general, a subjective conception of epistemic justification stipulates that a subject has justification for believing that p only if from his own perspective a situation obtains which makes likely the truth of p.² An objective conception of epistemic justification, on the other hand, stipulates that a subject has justification for believing that p only if the ground on the basis of which he forms the belief is such that it is (in some sense) reliably connected with the truth of p independently of whether or not it does so from the subject's own perspective.

² Later on, in Chapter 2, I will claim that on a subjective conception, a subject can also have

Now both the objective and subjective notions of epistemic justification as formulated face a number of problems. However, my present concern is not to address these problems but rather to try to understand Bonjour's objection to the phenomenological approach to the Basic Principle. How then should we understand Bonjour's claim that the phenomenal character of a perceptual experience has no bearing on the truth of the relevant proposition that we take the experience to provide justification for believing?

Suppose we accept some objective conception of epistemic justification. Perhaps we can then take Bonjour to be claiming that a perceptual experience with its characteristic phenomenal character is not as a matter of fact reliably connected with the truth of the relevant proposition. It should be immediately clear, however, that this is not the way we should understand Bonjour's challenge to the phenomenological approach. For this would be to say that a typical perceptual experience, in which it is to us as if a mind-independent object *x* has a certain property *F*, is not in fact reliably connected with the truth of the proposition that an *x* is *F*. But not even a radical sceptic about perceptual knowledge and justification would make this assumption. Nor, for that matter, would a proponent of the phenomenological approach want to be understood as saying that the Basic Principle can be explained by assuming merely that when one has a perceptual experience with a phenomenal character such that it is to one as if a mind-independent object *x* is *F*, such experiences are reliably connected with the truth of the proposition that an *x* is *F*. The problem here is, as far as the phenomenological approach is concerned, that this does not capture the claim that it is *in virtue* of a perceptual experience's phenomenal character that this experience provides one with epistemic justification for believing a certain proposition about one's environment. At issue, therefore, at

least insofar as we are assuming an objective conception of epistemic justification, is the question of what precisely accounts for the fact that a perceptual experience of a certain phenomenological sort is reliably connected with the truth of the relevant external world propositions.

Now in general, we can perhaps assume that if a type of perceptual experience is reliable with respect to the truth of the proposition p then it is because token experiences of this type are more likely than not caused by the fact that makes p true. But if this is all that we can assume, then at best what we can infer in epistemic terms is that the perceptual experience has the epistemic properties that it does in virtue of its being reliably caused by the fact that makes p true rather than in virtue of its intrinsic phenomenal properties. One way of understanding Bonjour's challenge, therefore, is as follows: if perceptual experiences provide us with any justification at all for believing propositions about our environment it is because these experiences are reliably connected with the truth of those propositions. However, there is nothing about the phenomenal character of perceptual experiences in virtue of which perceptual experiences are reliably connected with the truth of propositions about our environment or on the basis of which we could explain such reliable connections.

Given this understanding of Bonjour's challenge to phenomenological accounts of the Basic Principle, what options are there for a proponent of such an account? One option here is to turn to the metaphysics of perceptual experience or, more precisely, to the question of what constitutes an experience's phenomenal character. There are in fact various theories about the metaphysics of experience that are on offer. For example, following a sense-datum theorist, we might hold that a perceptual experience has the phenomenal character that it does as a result of the

subject's standing in some relation (e.g. of acquaintance or awareness) with a certain kind of mind-dependent object or entity. Though not popular these days, versions of the sense-datum theory still have their proponents.³ Alternatively, following an intentionalist theorist such as Pryor, we might think that a perceptual experience has the phenomenal character that it does in virtue of its being an intentional state with a certain representational content. Thus, when we have a perceptual experience which from the subjective point of view is such that a mind-independent object *x* is right before us having a certain property *F*, we are having an experience which represents an object *x* as being right before us and as having *F*.⁴ Finally, following a naive realist, we might think that a perceptual experience, at least when it is genuinely veridical and non-deceptive, has the phenomenal character that it does in virtue of the subject's standing in an irreducible relation of awareness to a mind-independent object and some of its properties. Thus, when we have an experience in which, subjectively speaking, a mind-independent object presents itself to us as having a property *F*, this, in the veridical non-deceptive case, is precisely because such an object and the instantiation by this object of some of its properties are constituents of the experience. Naive realism formulated as such, of course, cannot explain how certain kinds of hallucinations can have the phenomenal character that they do. It is natural, therefore, to combine this view with a disjunctivist conception of perceptual experience. On one such conception, a perceptual experience in which it appears to one as if an *x* is *F* can be either, in the veridical case, a perceptual experience whose

³ See, for example, (Jackson, 1977) and (Robinson, 1994)

⁴ We might want to distinguish here between what we might call strong-intentionalism and weak-intentionalism. Thus the strong intentionalist claims that a perceptual experiences has the phenomenal character that it does entirely as a matter of its having certain representational properties, whereas the weak intentionalist claims that the phenomenal character of a perceptual experience can be a matter of both its representational and non-representational properties. For defences of the former view see (Dretske, 1995), (Harman, 1990), and (Tye, 1995). For a defence of the latter view see (Peacocke, 1983)

phenomenal character is a matter of the subject's standing in some relation of awareness to a mind-independent object and some of its properties or, as in the hallucinatory case, a perceptual experience whose phenomenal character is a matter simply of this experience being subjectively indistinguishable⁵ from a veridical experience of an *x* being *F*.⁶

Now, given these three alternatives, it is not obviously clear how a sense-datum theorist or an intentionalist could explain how it is in virtue of its phenomenal character that a perceptual experience, in which it is to the subject as if a mind-independent object *x* is *F*, is objectively connected in the right sort of way to the truth of the proposition that an *x* is *F*. Matters, however, are different with the naive realist and the disjunctivist. Thus, at least in the veridical and non-deceptive case, these theorists say, as we have seen, that a perceptual experience has the phenomenal character that it does in virtue of the subject's standing in a certain relation to an object and some of its properties. If that is so, however, then a perceptual experience in which it is to the subject as if an *x* is *F*, will have as its constituents the object *x* and the instantiation by *x* of the property *F* and as such will entail (or if one prefers, be necessarily connected with) the truth of the proposition that an *x* is *F*. This form of naive realism, therefore, has the potential to make good on the claim that it is in virtue of its phenomenal character that a perceptual experience can provide us with epistemic justification for believing a proposition about the external world in the objective sense. Of course, more would need to be done to fill out a phenomenological account of the Basic Principle along these lines and, furthermore, some explanation would have to be given of how perceptual experiences can provide

⁵ To say that an experience *e* is subjectively indistinguishable from another experience *e*₁ is to say that one cannot know on the basis of introspection and reflection alone that *e* is distinct from *e*₁.

⁶ For discussion of this form of disjunctivism, see (Martin, *The Limits of Self-Awareness*, 2004) and (Martin, *On Being Alienated*, 2006). For other discussions on disjunctivism, see (McDowell, *Criteria, Defeasibility, and Knowledge*, 1982) and (Snowdon, 1980)

us with epistemic justification for believing propositions about the external world when these experiences are not veridical and deceptive. For moment, however, I want to leave this issue aside and turn to the question of how we can understand Bonjour's challenge if we presuppose a subjective conception of epistemic justification.

Recall that on one understanding of this conception, a subject has epistemic justification for believing a proposition *p* to the extent that, from his own subjective perspective, a situation obtains which entails or makes likely the truth of *p*. Earlier I suggested that, for the moment, we can think of the relevant subjective perspective as being constituted by some of the subject's beliefs. A subject, therefore, will have justification for believing a proposition *p* if he has an appropriate belief to the effect that a situation obtains which entails or makes it likely that *p* is true. Suppose then that we assume such a subjective conception of epistemic justification. If this is the case, we can perhaps understand Bonjour's challenge to consist in the claim that facts about the phenomenal character of a perceptual experience in which it is to the subject as if an object *x* is *F* do not provide the subject with any reasons for thinking that the proposition that an *x* is *F* is true.⁷ This means, in other words, that there is no reasonable or acceptable argument that the subject could provide, or any relevant piece of evidence that he could point to, that would show that facts about the phenomenal character of a perceptual experience in which it appears to one as if an *x* is *F* entail or make it likely that the proposition that an *x* is *F* is true. And it follows from this, so Bonjour might continue, that the subject cannot have an appropriate belief about the phenomenal character of the perceptual experience he is having

⁷ I take it that this is how Bonjour himself understands the challenge.

which would provide him with justification for believing the relevant proposition in the subjective sense.

Once again the question at this point is how a proponent of the phenomenological approach could respond. Without arguing for it here, I want to suggest that the proponent of this approach should concede to Bonjour that we cannot in fact provide a reasonable or otherwise acceptable argument, or point to any relevant piece of evidence, that would show that the phenomenal character of a perceptual experience in which it is to the subject as if an x is F entails or makes it likely that the proposition that an x is F is true. However, this concession to Bonjour need not entail the failure of the phenomenological approach, even on the assumption that the concept of epistemic justification should be understood in the subjective sense. For one option for the proponent of this approach to take here is to drop the assumption, which earlier we made provisionally, that a relevant subjective perspective on the basis of which a subject has any epistemic justification for believing a proposition p can only be constituted by some of the subject's beliefs. Instead, we might suggest that a relevant perspective can also be constituted by other mental states such as a subject's perceptual experiences (as well as his memories, intuitions, etc.) Furthermore, a proponent of the phenomenological approach might insist that it is precisely because of its phenomenal character that a perceptual experience can constitute a relevant subjective perspective on the basis of which a subject has some justification for believing a proposition p . Thus, it is because, in having a perceptual experience, it is to the subject as if a mind-independent object x has a property F that he occupies a perspective from which a situation obtains which entails or makes likely the truth of the proposition that an x is F .

Motivating a Phenomenological Approach

Again, the explanation of the Basic Principle about Perceptual Justification that was just sketched would need to be developed in more detail. But we have now seen, in a rough way, how a proponent of the phenomenological approach to the Basic Principle might attempt to address the challenge that Bonjour has presented to us.

At this point, however, we might ask: why should we take a phenomenological approach toward explaining the Basic Principle at all? One source of motivation for taking this approach, or so I would suggest, lies in the issues that arise from what we might call the new evil demon problem. To see what this problem amounts to, it will help to compare it to the familiar problem of Cartesian scepticism, which we might also call the old evil demon problem. Essentially, the latter is the problem of whether or not it is possible for us to have any perceptual justification for believing (and, ultimately, of whether or not it is possible for us to have any perceptual knowledge of) propositions about the external world. For the Cartesian skeptic, this problem about the possibility of perceptual justification, and ultimately of perceptual knowledge, arises from the assumption that we can have no justification for believing, and therefore no knowledge of, the proposition that we are not in fact victims of a certain kind of evil demon, who is responsible for our having the perceptual experiences that we in fact do and for our forming mostly false beliefs about our environment. Thus, the Cartesian skeptic assumes, for example, that we can have no justification for believing that we are the victim of an evil demon who is responsible for our having perceptual experiences as of there being mind-independent objects right before us and for our having beliefs about such objects

when in fact there are no such objects for us to be perceived. And it follows from this, so the skeptic contends, that we cannot have any perceptual justification for believing, and therefore cannot have any perceptual knowledge of, mundane propositions about the external world.

Now in contrast to this form of Cartesian scepticism, the new evil demon problem does not as such put into question the possibility of our having epistemically justified perceptual beliefs, or of our having any perceptual knowledge, about the external world. What it does put into question, however, is the viability of specific theories of perceptual justification and knowledge and, in particular, of theories that rely on the notion that a justified perceptual belief is one that is based on reliable grounds or produced by a reliable process. The problem here is nicely summarised by Ernest Sosa in the following passage:

What if twins of ours in another possible world were given mental lives just like ours down to the most minute detail of experience or thought, etc., though they were also totally in error about the nature of their surroundings, and their perceptual and inferential processes of belief acquisition accomplished very little except to sink them more and more deeply and systematically into error? Shall we say that we are justified in our beliefs while our twins are not? They are quite wrong in their beliefs, of course, but it seems somehow very implausible to suppose that they are unjustified. (Sosa, 1991, p. 132)

The intuition expressed in this passage, which I will refer to as the *new evil demon intuition*, is that a victim of an evil demon can have perceptually justified beliefs about his environment despite the fact that the perceptual grounds or processes on

the basis of which he forms these beliefs are highly unreliable. What seems to follow from this intuition is that the reliability of grounds or processes on the basis of which a subject forms his beliefs cannot be a necessary condition for these beliefs to be epistemically justified for him. More generally, what the new evil demon intuition seems to imply is that our perceptual beliefs cannot be justified in the objective sense of epistemic justification that we outlined earlier.

Now I take it that a phenomenological account of the Basic Principle of Perceptual Justification has to be motivated at least in part by its ability to provide an acceptable accommodation of the new evil demon intuition. And, on the surface, taking a phenomenological approach towards explaining how perceptual beliefs can be epistemically justified for us does seem promising in this respect. After all, what we and the subjects in the evil demon scenario seem to have in common is the fact that we both enjoy perceptual experiences that are of the same phenomenological kind.⁸ Thus if epistemic conclusions about how perceptual beliefs can be justified for us can be derived from facts about the phenomenal character of our experiences, then these conclusions are likely to apply not only to us but also to our counterparts in the evil demon scenario. Of course, none of this may turn out to be a conclusive reason for adopting a phenomenological account, even if such an account can be sustained. Thus, there may well be, for instance, other facts about the perceptual experiences of an evil demon victim or about the situation he finds himself in that equally apply to our experiences or the situation that we find ourselves in from which the relevant

⁸ This claim would need to be qualified if a disjunctivist conception of perceptual experience is correct. Thus, on this conception, it may be right to say that a veridical perception and a hallucination belong to the same phenomenological kind in the sense that both experiences are subjectively indistinguishable from the veridical experience. But, according to one particular form of disjunctivism a veridical perception and a hallucination essentially also belong to two different phenomenological kinds since the phenomenal character of the former, but not the latter, is constituted in by a relation to a mind-independent object

epistemic conclusions about how our perceptual beliefs are justified for us can be derived.

However, the new evil demon intuition does, I believe, provide a different motivation for adopting a phenomenological approach to the Basic Principle. For one point to note is that the presentation of the evil demon scenario from which we elicit the intuitive judgments about when a perceptual belief can be epistemically justified for a subject is based in large part on the assumption that the victims of the evil demon have perceptual experiences that are, phenomenologically speaking, just like the perceptual experiences that we enjoy.⁹ In other words, it is precisely because of the assumption that the victims of the evil demon have perceptual experiences that are phenomenologically just like our perceptual experiences that we form the intuitive judgment that they, like us, have perceptual beliefs about their environment that are epistemically justified. This would seem to suggest that insofar as the new evil demon intuition is true and justified for us that it is because of the phenomenal character of their perceptual experiences that the victims of the new evil demon have perceptual beliefs that are epistemically justified for them. And if it is in virtue of the phenomenal character of their perceptual experiences that these victims have perceptual beliefs that are justified for them, it may seem reasonable to assume that the same conclusion applies to our perceptual beliefs as well.

Outline of the Thesis

⁹ Again, see n. 8

The structure of this thesis is as follows. In Chapter 1, I will consider the question of whether it is possible to provide an explanation of the Basic Principle in terms of the objective conception of epistemic justification. As we have seen such an explanation is essentially confronted with the new evil demon problem. I will thus consider potential accounts of the Basic Principle in terms of the objective conception of epistemic justification that specifically attempt to accommodate the new evil demon intuition. I will be concerned, in particular, with the accounts given by Michael Bergman and Tyler Burge and argue that neither of these provides a successful account of perceptual justification.

In Chapter 2, I will turn to potential accounts of the Basic Principle in terms of the subjective conception of epistemic justification. I first consider and reject as a possible account Richard Foley's conception of epistemic rationality. I then turn to John McDowell's conception of perceptual justification, which I show can be taken to provide an explanation of the Basic Principle in terms of both the objective and subjective conception of epistemic justification. While this conception is ultimately unsatisfying, I suggest that it can illuminate a conception of perceptual justification in terms of the subjective conception of epistemic justification that can successfully accommodate the new evil demon intuition. Crucially, this account (like, in fact, the McDowellian one) appeals to the phenomenal character of perceptual experience.

In Chapter 3, I consider potential connections between the metaphysics and epistemology of experience. In particular, I consider the potential epistemological implications of a disjunctivist conception of perceptual experience. Since on such a conception, the phenomenal character of a veridical perceptual experience is constituted differently than the phenomenal character of a hallucination, a disjunctivist account, insofar as perceptual justification is a matter of the phenomenal

character of an experience at all, leaves potential room for the claim that in the veridical case, perceptual experiences can provide one with a kind of justification that is not available in the hallucinatory case. I analyse some arguments in McDowell's writings that can be taken to support the claim that veridical perceptual experiences can, and should be taken to be able to, provide a kind of justification unavailable in the hallucinatory case. In turn, these arguments, if sound, might ultimately be taken to support the claim that veridical experiences can provide a kind of justification unavailable in the hallucinatory case because of the way the phenomenal character of the veridical experiences is constituted. I argue, however, that McDowell's arguments fail.

In the final Chapter 4, I turn away from considerations about the phenomenal character of perceptual experience and turn instead to the question of whether or not perceptual experiences can provide us with a kind of immediate justification, that is, with justification that does not depend on our having independent justification for further beliefs, for example, for the belief that perceptual experience are generally reliable. I address this question in terms of the problem of easy knowledge as it is discussed by philosophers like Cohen. I argue that the problems Cohen raises need not compel us into thinking that perceptual justification cannot be immediate. As such, at least as far as the issues concerning easy knowledge are concerned, the Basic Principle can be taken to be made true by the fact that perceptual experiences can provide us with epistemic justification that does not depend on our having independent justification for further beliefs.

Chapter 1

Introduction

What we seek in this thesis is an explanation for the truth of the Basic Principle about Perceptual Justification. This principle, in its doxastic form, states that:

If S has a perceptual experience as of a mind-independent object x being F and forms the belief that an x is F on the basis of this experience, then S's belief that an x is F is prima facie justified for S.

The notion of justification mentioned in this principle is epistemic and therefore an evaluative notion, which can be understood in the following sense: to say that a belief that p is justified is to say, among other things, that it is appropriate in some sense with respect to the cognitive aim of believing p only if p is true. On an objective conception of justification, this may mean that a belief that p is justified only if it is based on grounds that in fact reliably indicate the truth of p or, alternatively, only if it is produced by a belief-forming process which in fact reliably produces true beliefs. On a subjective conception of justification, it may mean that a belief that p is justified only if the subject has an appropriate belief to the effect that conditions obtain that make it more likely than not that p is true. In this chapter, I address the question of whether it is possible to provide an explanation of the Basic Principle of Perceptual Justification in terms of the objective conception of epistemic

justification. I will refer to such explanations of the Basic Principle as *objectivist* explanations.

One such explanation is what we might call a standard form of reliabilism about perceptual justification. This form of reliabilism holds that the Basic Principle is true insofar as a perceptual experience of an *x* being *F* is a reliable indication that an *x* is *F* or insofar as the belief that an *x* is *F* when formed on the basis of such a perceptual experience is formed by a belief-forming process that reliably produces true beliefs to effect that an *x* is *F*. Standard forms of reliabilist accounts of epistemic justification, however, face several well-known problems. One of these is the problem of clairvoyance as it is presented by Bonjour.¹⁰ This problem concerns a subject who is capable of forming reliably true beliefs on the basis of a clairvoyant power. Bonjour considers various cases of a clairvoyant subject each differing in terms of the subject's having or lacking evidence against having such powers, etc. The appeal to these cases is supposed to show that there can be instances in which a subject forms beliefs on the basis of highly reliable grounds or processes but lacks any justification for having these beliefs. If this really is the consequence of cases involving such a clairvoyant subject, it will mean that reliability of grounds or processes cannot be a sufficient condition for having justified beliefs formed on the basis of these grounds or processes.

Another problem that is faced by standard forms of reliabilism is the new evil demon problem introduced earlier. As we have seen, this problem seems to show that reliability of grounds or processes cannot be a necessary condition for having justified beliefs formed on the basis of these grounds and processes. It is this

¹⁰ (Bonjour, *The Structure of Empirical Knowledge*, 1985)

problem, and the way it confronts objectivist explanations of the Basic Principle in general, that I want to address in this chapter.

If we accept the intuition that a victim of an evil demon is capable of having justified perceptual beliefs and therefore reject any form of standard reliabilism about perceptual justification, there seem to be two possible options we can take in order to explain the Basic Principle while retaining at least in part an objectivist conception of epistemic justification. The first of these options appeals to the idea that there are two senses in which a subject's beliefs can be epistemically justified for him. This is a proposal that has been made by Alvin Goldman.¹¹ In the first sense, which Goldman refers to as the 'strong sense', a subject's beliefs are justified insofar as these beliefs are formed on the basis of reliable grounds or a reliable process. This is the sense in which a subject's perceptual beliefs are generally justified in friendly worlds or environments of the kind that we take ourselves to inhabit. In the second sense, which Goldman refers to as the 'weak sense', a subject's beliefs are justified insofar as he is epistemically blameless in holding these beliefs. This is the sense in which a subject's perceptual beliefs are generally justified in unfriendly worlds or environments of the kind inhabited by victims of an evil demon. According to Goldman, a subject S is *merely* justified in holding a belief that p in this second, weak sense when the following conditions are met:

- (1) The method M by which the belief is produced is unreliable.
- (2) S does not believe that M is unreliable
- (3) S neither possesses, nor has available to him/her, a reliable way of telling that M is unreliable.

¹¹ (Goldman, Strong and Weak Justification, 1988)

- (4) There is no process or method S believes to be reliable which, if used, would lead S to believe that M is unreliable. (ibid, 56)

The idea is thus roughly that a subject is justified in (merely) the weak sense insofar as his beliefs are unreliably formed but he cannot be reasonably blamed for doing so or for not respecting any countervailing evidence that there might be. However, in the remainder of this thesis, I will not be addressing in any detail the proposal that any form of epistemic justification can be accounted for in terms of doxastic blamelessness. The reason, briefly, is that the notion of doxastic blamelessness does not seem to be an epistemic notion at all since it does not seem to be conceptually connected in a relevant way with the notion of truth. Thus, suppose, for example, that I suffer from a particular brain lesion that causes me form the belief in the proposition that I am constantly being stalked and that prevents me from considering all the countervailing evidence that in fact exists against this proposition. I am certainly blameless in holding this belief but unless it is more likely than not to be true or unless I have some good reasons for thinking that I am constantly being stalked it is far from clear why this belief should be in any way epistemically justified for me.¹²

The focus in the remainder of this chapter will therefore be with evaluating versions of the second option mentioned earlier. Roughly, this response to the new evil demon problem on the part of a proponent of an objectivist account of perceptual justification consists of saying that while reliability in the subject's own world or environment of the grounds or processes on the basis of which he forms his beliefs may not be a necessary condition for those beliefs to be epistemically

¹² This is not to say, however, that we and the new evil demon victims cannot be justified in two different senses. I discuss a potential two sense account in Chapter 3

justified for the subject, it is a necessary condition for these beliefs to be justified that the grounds or processes on which they are based are reliable relative to *some* world or environment. As I will be showing, this kind of objectivist response to the new evil demon reasoning faces one particular fundamental problem. We can get an initial handle on this problem by considering the following example of an objectivist account. Suppose we were to say that a victim of an evil demon had justified perceptual beliefs because his perceptual experiences provided him with grounds that, while unreliable in his own world or environment, are reliable in our, the actual world.¹³ This objectivist proposal would be immediately confronted with the following question: why should it be, objectively speaking, epistemically appropriate for the victim of an evil demon to form beliefs about his environment on the basis of grounds or processes that are reliable relative to some world or environment but highly unreliable relative to the world or environment that he finds himself in? Call this the *epistemic relevance question*. I shall argue that the prominent objectivist accounts of perceptual justification that can be found in the literature which relativise reliability of perceptual grounds or perceptual belief forming processes face real difficulties in answering, and ultimately fail to provide satisfying answers to, this question.

I will approach this issue as follows. In the next section, I discuss how the epistemic relevance question arises for standard forms of reliabilism independently of the new evil demon problem and how a standard reliabilist might satisfactorily address this question. I will then consider two objectivist accounts of perceptual justification that can be taken as attempting to solve the new evil demon problem by relativising reliability of perceptual grounds or perceptual belief forming processes

¹³ See (Goldman, *Epistemology and Cognition*, 1986) for a proposal of this kind. Goldman does not actually endorse this proposal

to worlds or environments other than the world or environment inhabited by the victim of the new evil demon. In particular, I will be looking at the objectivist accounts of perceptual justification that are offered by, or can be derived from the work of, Michael Bergman and Tyler Burge.

Global Reliabilism

Reliabilism in its standard form faces questions concerning some (variations of certain) well-known thought experiments other than the new evil demon scenario. These questions are raised by a recent discussion about reliabilism by David Henderson and Terry Horgan. Henderson and Horgan appeal to a thought experiment that goes like this:

Suppose that Athena and Fortuna are driving...through a county in which there happen to be numerous extremely realistic-looking fake barns within view of the highway – although neither of them has any inkling of this fact or any reason to suspect it. As it happens, in this local area all the real barns are yellows, and none of the fake barns or any other buildings are yellow. Again, they have no information to this effect. As they drive past a saliently presented yellow building, Athena, who has had reasonable experience with barns, gets a clear look at it, and on the basis of its barn-like visual experience, she judges it to be a barn.

Fortuna only gets a very brief glimpse of the building. She saw her first barn just yesterday, elsewhere, and it happened to be yellow. She judges, on

the basis of the briefly glimpsed building's yellow color, that it is a barn – even though she did not get a good look at it, and was thus unable to discern any features that are generally distinctive of barns as opposed to other kinds of buildings. It's not that she has a *general* belief that all and only yellow buildings are barns, or that all barns are yellow...Also, it's not that she has a *general* tendency to inductively extrapolate from old cases to new ones in a hastily generalizing way, a tendency she might otherwise be exhibiting here. Rather, it just happens that in the present circumstances, a psychological process is present within her that takes as input both the brief glimpse of a yellow building and the yellow-barn memory, and generates as output the barn-belief about the briefly glimpsed object.

So Athena and Fortuna each form the belief that the building is a barn. And indeed it is. (Henderson & Horgan, 2007, pp. 103-4)

According to Henderson and Horgan, the intuitive verdict about Athena in this case is that while she fails to know that the building is a barn she nonetheless has a justified belief that it is despite the fact that the process on the basis of which she forms that belief – which takes a perceptual experiences of barn-like structures as an input and produces as an output the belief that the building is barn – are unreliable in the local environment through which she and Fortuna are currently driving. The intuitive verdict about Fortuna, on the other hand, is that she neither knows nor has a justified belief that the building is a barn despite the fact that the process on the basis of which she forms the belief – which takes a perceptual glimpse of a yellow building as an input and produces as an output the belief that the building is a barn – is a reliable in the given local environment. Still, Henderson and Horgan maintain

that neither of these two intuitive verdicts concerning the beliefs of Athena and Fortune entail that reliabilism about epistemic justification is essentially wrong. Rather what these verdicts may indicate to us is that what matters to one's having an epistemically justified belief is that this belief is formed on the basis of grounds or processes that are *globally* rather than merely locally reliable.

A subject's global environment can be understood as a set consisting of various local environments that the subject may find himself in. Globally reliable belief-forming processes are thus ones that have "a tendency to produce (mostly) true beliefs...with respect to the wide reference class comprising the potential local environments to which an agent might be exposed (or which the agent might inhabit) within that agent's global environment." (ibid, 105) On the emerging *global reliabilist* view, therefore, Athena's belief that the building is a barn is justified because it is based on a process that is globally reliable – after all, most barn-looking structures across the global environment are real barns – whereas Fortune's belief is unjustified because it is based on a process that is merely locally reliable.

What should we make of this proposal? One question we might ask here is whether there is any motivation for thinking of epistemic justification as a matter of global reliability in Henderson and Horgan's sense independently of whether doing so renders correct the intuitive verdicts concerning the epistemic status of Athena's and Fortune's beliefs in the above case. Henderson and Horgan certainly think that such independent motivation can be provided. This is made clear in the following passage:

One can begin to appreciate why epistemic safety – and the relative robustness of reliability – of belief-fixing processes would be

epistemologically significant when one reflects on a prominent and pervasive characteristic of epistemic life: one's epistemic endeavour must be undertaken in the face of uncertainty. One's epistemic chores must be managed while possessing only a fallible understanding regarding one's global and local environment. Such epistemic uncertainty (or fallibility) regarding one's environment is paralleled by, or mirrored in, an uncertainty (or fallibility) regarding what processes will work in one's environment. In view of the uncertainty characteristic of the epistemic situation, consider two alternatives...On the one hand, one could employ a process that is reliable with respect to the wide reference class comprising the potential local environments two which one might be exposed (or which one might inhabit) within one's global environment...[Such a process] is relatively safe insofar as its reliability does not depend heavily upon certain unusual or atypical features that highly specific to the particular circumstance or environment which the possessor of the disposition might happen to occupy; its reliability then does not obtain only relative to a narrow reference class of environments in which those particular features happen to be present. Alternatively, one might employ processes that would be reliable only given certain unusual or atypical environment features – only given features that are highly specific to particular circumstances or environments within the global environment that one happens to occupy. If, in employing such a process one employs a reliable process, it is a merely locally reliable process. In light of the uncertainty that is a pervasive fact of epistemic life, it is clear what one should make of this general abstract choice. The safety afforded by globally reliable processes would be rationally desirable in preference to the risk one

runs when using a merely locally reliable process. (Henderson & Horgan, 2007, pp. 107-8)

The argument that Henderson and Horgan present here seems to be this: we should think of an epistemically justified belief as one that is produced by a globally reliable process because choosing such a process would be rationally preferable to choosing a process that is merely locally reliable.

But now the question is: ‘rationally preferable’ in what sense? In Henderson and Horgan’s sense, the rationally preferable choice here arises from a particular set of constraints. First a choice has to be made between two competing belief-forming processes, namely a globally reliable process and a merely locally reliable one. Second, the choice has to be made under conditions in which the chooser lacks knowledge of relevant features of the environment that he finds himself in and knowledge of what belief-forming processes will work in that environment. Third, though this is not explicitly stated in the above passage, we can take it that the choice has to be made with the aim of forming a true belief about what obtains in the given environment. Thus the sense in which a belief-forming process ‘works’ or does not ‘work’ in the given environment is that of whether this process will or will not produce a belief that is more likely than not to be true.

Now given these constraints, we can concede that it straightforwardly follows that the globally reliable process is by far the rationally preferable choice. For given one’s antecedent knowledge about the choices one has and about how these might work in the given environment, in choosing the globally reliable process one minimizes the risk of forming a false belief. Hence it is the globally reliable process rather than the merely locally reliable one that can provide one with justification for

believing the relevant proposition in one's current local environment – even if this process turns out to be unreliable in that environment – because this process ensures a level of epistemic safety that cannot be secured by the merely locally reliable one. Of course, it should be said here, that epistemic subjects under ordinary circumstances do not actually make the kind of choice that is envisaged in the above passage. But as Henderson and Horgan make clear, this kind of choice is merely an abstract one. Thus we can take the conditions under which this kind of choice is made as representing a certain rational perspective which fixes for us an objective conception of epistemic justification; that is, a conception of when beliefs are or are not, objectively speaking, appropriate from an epistemic point of view. From this perspective, a belief based on globally reliable processes is epistemically appropriate whereas a belief based on a merely locally reliable process is not.

How good is the argument that Henderson and Horgan provide in support of the claim that we should think of epistemic justification as a matter of the global reliability of a given belief-forming process? One further question that we might have at this point is why we should think that it is the particular perspective of a rational subject that Henderson and Horgan rely on that should fix for us an objective conception of epistemic justification. Certainly, we can think of relevant rational choices being made under very different constraints. For example, we could suppose that a rational subject – faced with the two choices for belief-forming processes as Henderson and Horgan's rational subject – does in fact have prior knowledge of whether or not the two processes will work in the local environment in which the relevant belief is to be formed. Of course, from this perspective, it will be, objectively speaking, entirely appropriate to rely on the locally reliable process rather than the globally reliable one. After all, the locally reliable process will in fact

make it more likely than not, in the given case, that the belief to be formed is true. So why should this rational perspective rather than the one outlined by Henderson and Horgan not fix for us an objective conception of epistemic justification?

At this point, I want to suggest that rather than pursuing this line of questioning, that the issue here is far more straightforward than Henderson and Horgan suggest. The question, to repeat, is why it should be epistemically appropriate, objectively speaking, for us to rely on globally reliable processes in forming our beliefs? The answer to this question seems to be just this: we should do so because we do, as a matter of fact, conduct our epistemic affairs within a global environment and are not, therefore, merely confined to just a particular local environment. It follows from this that relying on globally reliable processes will ensure that most of the beliefs that we form in conducting our epistemic affairs will turn out to be mostly true. And given that it is surely a proper epistemic goal to have beliefs that are all, or mostly, true, this makes reliance on globally reliable processes epistemically appropriate from an objective point of view. It also follows from the fact that we do conduct our epistemic affairs within a global environment that relying on globally reliable processes is epistemically more appropriate than relying on a merely locally reliable process because the latter unlike the former cannot ensure that most of our beliefs turn out to be true in a global context.

At this point, we might ask, however, if we do conduct our epistemic affairs within a global environment, why – objectively speaking – we should rely on a globally reliable process rather than relying on whatever process might be locally reliable in the environment that we happen to find ourselves in. For objectively speaking, reliance on whatever process is reliable in the local environment we happen to occupy at any given moment may satisfy the epistemic goal of having

mostly true beliefs in a global context even better than mere reliance on a globally reliable process.

In response to this question, I think that the reliabilist can concede that relying on whatever process is reliable in the local environment we happen to occupy may be even more objectively appropriate than merely relying on a globally reliable process but still insist that doing the latter is nonetheless epistemically appropriate to a sufficiently high degree such that beliefs produced by such a globally reliable process count as epistemically justified. In other words, relying on processes that are globally reliable does a sufficiently good job of fulfilling the aim of having true beliefs and avoiding false ones within a global context in way that reliance on such processes can result in genuinely justified beliefs.

Suppose then that we ask the epistemic relevance question in relation to Henderson and Horgan's Athena: why should she rely on a process that is globally reliable but unreliable in the local environment she finds herself in? The answer to this question is that, even though the globally reliable process is more likely than not to produce a false belief in the local environment that Athena finds herself in, it nonetheless ensures that most of the beliefs that she forms on the basis of this process will come out to be true rather than false. So in the absence of her having reasons for thinking that she is in a local environment in which her processes are unreliable, it is epistemically appropriate for her, to a degree sufficient to yield epistemic justification, to rely on her barn-like visual experiences.

The kind of thought experiment that Henderson and Horgan run therefore provides us with some motivation for adopting a form of global reliabilism. In the final analysis, however, this sort of reliabilism, of course, is still unacceptable because it fails to address the new evil demon problem. Thus we can imagine a

scenario in which a victim of an evil demon enjoys perceptual experiences that from the inside are just like our perceptual experiences but inhabits a global environment in which these perceptual experiences are highly unreliable. As we have seen, the intuitive judgment about such victims is that they can have justified perceptual beliefs which means that the reliability of their perceptual experiences relative their global environment cannot be a necessary condition for their having justified perceptual beliefs.

So we are still seeking an objectivist account of the Basic Principle that can accommodate the new evil demon intuition. As I have indicated, the objectivist approaches that I will be considering in more detail will insist that being epistemically justified requires having beliefs that are based on grounds or processes that are reliable relative to some environment, even if they are not reliable relative to the environment that a subject, like the evil demon victim, might find himself in. However, just like the global reliabilist, the proponent of such an approach will have to provide some principled explanation for why it should be epistemically appropriate, from an objective point of view, for a subject to rely on processes or grounds that are not reliable in the environment he finds himself in.

In next section, I will consider the accounts of the justification of perceptual beliefs that are suggested by, or can be derived from, the work of Bergman on the one hand and Burge on the other. I will first provide an outline of each view before evaluating their viability with respect to the epistemic relevance question.

Bergman's Proper Function Account

In a recent book, Bergman develops an objectivist conception of an epistemically justified belief as one that is produced by a properly functioning cognitive faculty. The most detailed development of epistemic notions in terms of proper function is provided by Alvin Plantinga.¹⁴ Plantinga is concerned with giving an account of what he calls *warrant* which is a property that a belief can have and which turns a true belief into knowledge. Instead of analysing Plantinga's account of warrant, however, I will focus on the recent account provided by Bergman. Like Plantinga, Bergman appeals to the notion of proper function in developing an epistemological theory. But unlike Plantinga, Bergman is concerned with the concept of epistemic justification. This concept, as we have been understanding it, is a concept of a property that makes a belief appropriate from an epistemic point of view but is not one which essentially turns a true belief into knowledge. Even more importantly for our current purposes, Bergman's account of epistemic justification in terms of proper function is intended in part to account for the problem presented to us by the new evil demon scenario. If viable, this account could provide us with an objectivist conception of epistemic justification as well as an explanation of the Basic Principle of Perceptual Justification consistent with the new evil demon intuition.

As a general analysis of what an epistemically justified belief amounts to, Bergman offers the following:

S's belief B is justified iff (i) S does not take B to be defeated and (ii) the cognitive faculties producing B are (a) functioning properly, (b) truth-aimed and (c) reliable in the environment for which they are 'designed' (Bergman, 2006, p. 133)

¹⁴ (Plantinga, 1993)

If we apply this general analysis of an epistemically justified belief to the Basic Principle of Perceptual Justification in its doxastic form, we thus say that the Basic Principle is true insofar as S's forming the belief that something is x on the basis of having an experience as of an x being F is the result of the proper functioning of a cognitive faculty that is truth-aimed and reliable in the environment for which this faculty is 'designed'. The key elements of this analysis are the conditions listed under (ii). These conditions appeal to the notion of a cognitive faculty that is functioning properly, aimed at producing true beliefs and thus 'designed' to be reliable under certain conditions. How should we understand these notions? As Plantinga notes, the notions employed here are most naturally applied to artefacts and devices and the purposeful, conscious beings who design them. (Plantinga 1993, 195) This immediately raises the question of how we can apply these notions to our cognitive faculties without having to assume that these are artefacts or devices designed by some purposeful and conscious being. For if we don't take our cognitive faculties to be the product of a purposeful, conscious designer, in what sense should we understand the notion that these faculties are functioning properly, aimed at truth and therefore 'designed' for a specific purpose? Ultimately, we might think that such notions are inapplicable to our cognitive faculties unless they can be given some naturalistic explanation. However, I will leave these issues aside here. Furthermore, to simplify our discussion, I will use these notions as we most naturally understand them, namely as applying to artefacts and devices and as presupposing a purposeful, conscious designer.

Why then should we take a justified belief as a belief that meets the conditions under (ii)? Consider first condition (ii)(a). Bergman's explanation of why this

condition holds for justified beliefs runs allowing the following lines. Take the case of a perceptually justified belief. Generally, we might say that a belief that an x is F when formed on the basis of an experience as of an x being F is justified insofar as that belief is an ‘epistemically fitting doxastic response’ to the perceptual experience. An account of what makes a doxastic response a fitting response to a perceptual experience may accept the following three conditions:

Nonreliability: the fittingness of doxastic response B to evidence E is not contingent upon E’s being a reliable indicator of B’s truth [in the subject’s environment]

Objectivity: the fittingness of a doxastic response B to evidence E is objective fittingness (in the sense that fittingness from the subject’s perspective isn’t sufficient for it).

Necessity: the fittingness of doxastic response B to evidence E is an essential property of that response to that evidence. (Bergman 2006, p. 112)

For reasons I won’t go into here, Bergman argues at length that an account of what constitutes a fitting doxastic response to a piece of evidence should drop Necessity. Applied to our perceptual case from before this means that while for a subject like us the belief that an x is F may, in the absence of any countervailing evidence, be a fitting doxastic response to his having a perceptual experience of an x being F, this need not be the case for subjects who are unlike us in some epistemically relevant respect. So what then makes a belief a fitting doxastic response to a perceptual

experience if what does so is not a necessary but contingent property of that response? According to Bergman, what makes a belief a fitting doxastic response to a perceptual experience, is that this belief is produced by a properly functioning cognitive faculty. Hence condition (ii)(a) is a necessary condition for having a justified belief. Bergman's suggestion may thus be taken to mean that our cognitive faculties are designed in such a way that when we have a perceptual experience of an x being F, under normal circumstances, we respond by forming the belief that an x is F. Other species, however, may have cognitive faculties that are designed to form different kinds of beliefs on the basis of the same kind of perceptual experience. This does not preclude that their beliefs are doxastically fitting responses to a perceptual experience of that kind as long as further conditions are met.

These conditions are (ii) (b) and (ii) (c). (ii) (b) is required to make the response produced by a properly functioning cognitive faculty an epistemically fitting as opposed to say, a pragmatically fitting, response. The following example illustrates this point:

...suppose...that the particular faculty producing the belief we are considering is intended by its creator not to produce *true* beliefs but, rather, to produce beliefs that will minimize psychological trauma (even if that involves regularly producing false beliefs). Then it seems that beliefs being produced by such a cognitive faculty won't be *epistemically* fitting responses to the input to the subject's belief-forming system though they may be appropriate in some other sense. This would be a case of a belief that isn't justified (since it isn't an epistemically fitting doxastic response) even though it is produced by a properly functioning cognitive faculty. (Bergman, p. 135; his emphasis)

A similar point applies to condition (ii)(c). That is, unless the cognitive faculty is reliable in the way it was designed to be, it will not produce epistemically fitting responses to inputs in a subject's belief-forming system and will therefore be incapable of producing epistemically justified beliefs. I take it that this condition is meant to satisfy the general objectivist requirement that a belief-forming process yields epistemically justified beliefs only if it is in some sense sufficiently truth-conducive independently of its being truth-conducive from the subject's perspective. Here is Bergman's example in support of this claim:

...consider a creature designed by one of Hume's infant deities. And suppose that, although this incompetent creator was trying to make a believer with reliable faculties, it instead created one whose faculties produce mostly false beliefs when placed in the environment in which it was intended by its creator to produce true beliefs. For example, suppose this infant deity intentionally created Ric in such a way that...his natural unlearned doxastic response to an [experience] ME2 is to form [belief] B1. But, contrary to what this bumbling creator had hoped, this belief (like most other belief Ric forms) is false when produced in the intended environment by faculties functioning as they were designed to function. Is Ric's belief B1 an epistemically fitting response to ME2? It seems not...Despite the failed design attempt, the belief in question may be, in some sense, the output of properly functioning cognitive faculties...But because the plan, when implemented as intended, didn't work, the belief doesn't seem to be an epistemically fitting response to the subject's evidence. (ibid, pp. 135-6).

Of course, the stipulation that the cognitive faculty only has to be reliable in the environment in which it is intended to be reliable, rather than the environment in which the subject of the faculty may happen to find himself in, is such that it can accommodate the new evil demon intuition. Thus on Bergman's analysis, the victim of the evil demon can have justified perceptual beliefs, even if the cognitive faculties producing these beliefs are highly unreliable in his environment, provided these faculties are functioning properly, aimed at producing true beliefs and reliable in the environment they are intended to be reliable in.

Is this a satisfactory accommodation of the new evil demon intuition? Before answering this question, I want to first look at an alternative objectivist account of perceptual justification.

Burge's Anti-Individualist Account

The other account is the view about perceptual entitlement that has been developed by Tyler Burge. Before analyzing this view, it will be helpful to first clarify Burge's use of certain terms. The term 'entitlement', as Burge uses it, refers to a sub-species of what he calls 'epistemic warrant'. The latter is used to refer to what I have been calling justification here rather than to what Plantinga refers to by his use of the term 'warrant' which is a property of belief that turns a true belief into knowledge. According to Burge, an 'entitlement is *epistemically externalist* inasmuch as it is warrant that need not be fully conceptually accessible, even on reflection, to the warranted individual. The individual need not have the concepts necessary to think

the propositional content that formulates the warrant.” (Burge, 2003) The other subspecies of warrant is what Burge refers to as ‘justification’, which is “warrant by reason that is conceptually accessible on reflection to the warranted individual” (ibid, p. 505) In explicating a conception of perceptual entitlement, Burge seeks an understanding of how beliefs about one’s environment can be warranted by one’s perceptual experiences where this warrant is not derived from a subject’s possession of reasons from which the relevant propositions about the environment can be inferred. Perceptual entitlements for beliefs about the environment are therefore essentially non-inferential. In what follows, I shall speak of Burge’s view on perceptual justification rather than entitlement. In doing so, I will use the term ‘justification’ in the way that I have been doing so far. This use is to be distinguished from Burge’s use of the same term.

The key element of Burge’s view on perceptual justification is that it derives in part from the truth of *anti-individualism* about representational states, namely “...the principle that certain representational states are dependent on their natures, and for the individuation of what representational contents they have, on certain relations between the individual or relevant representational systems of the individual – and certain aspects of the environment that is represented.” (ibid, 505) To Burge, perceptual experiences are just one type of representational state whose nature is individuated in this way. Thus he states:

A condition on particular perceptual representational states’ having the content that they have is that there have been both causal-formative interactions (which are not in themselves representational) and representationally successful interactions between instances of types of relevant perceptual referents and

aspects of the individual's perceptual system (in either the individual's history, or in evolution of the system in his evolutionary ancestors, or in some other way). (ibid, p. 531)

On this view, a perceptual experience's representing an *x* as being *F* can only be a result of actual causal interactions between instances of *x* and instances of the property *F* on the one hand and the perceptual system of the subject whose perceptual experience it is on the other. In connection with this form of anti-individualism about perceptual experiences, Burge introduces the notion of a *normal environment*. This is the environment "by reference to which the perceptual content of the perceptual state is explained and established." (ibid, p. 532) In other words, it is the environment in which causal interactions occur between features of that environment and the subject's perceptual representational system, causal interactions which make the subject's perceptual experiences have the representational contents that they do.

The crucial epistemic move that Burge makes then comes at this point. Suppose that the causal interactions between features of the subject's normal environment and his perceptual representational system are the result of lawlike connections in such way that the representational content of these states are reliably veridical in the normal environment. One upshot of this supposition is that the reliable causal connections between features of the subject's normal environment and his perceptual representational system are constitutively connected with the nature of these states; in other words the reliable causal connections make the perceptual states what they are, namely states with certain representational contents. Crucially for Burge, when perceptual experiences and their representational contents are constituted in this way

by reliable causal connections between the subject's perceptual representational system and the environmental features that his experiences represent, the perceptual experiences become epistemically efficacious. For essentially, according to Burge, a perceptual experience with a certain representational content of an x being F provides one with justification for believing the corresponding proposition that an x is F only if perceptual experiences of this type are reliably veridical in the subject's normal environment and only if the reliable veridicality of perceptual experiences of this type is constitutively connected to these experiences having the nature that they do.

At this point, I want to briefly compare my analysis so far of Burge's view of perceptual justification so far with Anthony Brueckner's interpretation of the same view. Brueckner summarises Burge's view of perceptual justification as follows:

- I. Anti-individualism holds for perceptual content. (Assumption)
- II. An individual's perceptual state types are reliably veridical in his perceptual system's normal environment. (I)
- III. This reliable veridicality is explained by the nature of the perceptual states. (I)
- IV. If (i) an individual's perceptual state types are reliably veridical in his perceptual system's normal environment, and (ii) this reliable veridicality is explained by the nature of the perceptual states, then the individual is entitled to hold beliefs appropriately based upon those states (Assumption).
- V. Perceptual entitlements exist. ((II), (III), (IV))¹⁵

¹⁵ (Brueckner, 2007, p. 162)

Brueckner interprets Burge as saying that (I) entails (II). On this interpretation, we would take Burge to say that we can validly infer a priori from the assumption that perceptual experiences have contents that are anti-individualistically individuated the proposition that these experiences are reliably veridical in the subject's normal environment. However, as far as I can tell, Burge nowhere makes an assumption along these lines. In the first instance, in line with his acceptance of anti-individualism about perceptual representational content, he simply assumes that perceptual experiences have the contents that they do as a result of actual causal interactions between features of the subject's home environment and his perceptual representational system. Burge does not seem to hold in addition to this that perceptual experiences have the representational contents that they do only if the relevant causal interactions that individuate these contents are also reliable. Rather his view seems to be that if perceptual experiences have representational contents that are anti-individualistically individuated and if, as a matter of contingent fact, these experiences have the contents that they do as a result of reliable causal connections between the relevant features of the subject's normal environment and his perceptual representational system, then these perceptual experiences can provide the subject with epistemic justification (or entitlement) for forming beliefs about his surroundings.

In summary, my analysis of Burge's view on perceptual justification, unlike that of Brueckner's analysis, is therefore as follows

- I. Anti-individualism holds for perceptual content. (Assumption)
- II. An individual's perceptual state types are reliably veridical in his perceptual system's normal environment (Assumption).

- III. The nature of the individual's perceptual states is explained and constituted by this reliable veridicality. (I, II)
- IV. If (i) an individual's perceptual state types are reliably veridical in his perceptual system's normal environment, and (ii) the nature of the perceptual states is explained and constituted by this reliable veridicality, then the individual is justified (or entitled) to rely up on his perceptual states, and he is justified (or entitled) to hold beliefs appropriately based upon these states. (Assumption)
- V. Perceptual justification (entitlement) exists. ((II), (III), (IV))

As we have seen, it is the last premise in this argument that is the key thought in Burge's conception of perceptual justification. To him, it is absolutely crucial that mere reliable veridicality of perceptual experiences does not by itself render these experiences epistemically efficacious. Thus he claims: "Reliable connections to the world that are accidental relative to the conditions that individuate the individual's perceptual states and competencies contribute nothing to empirical epistemic entitlement." (Burge, 2003, p. 534) This means that accidental connections – that is, connections that do not contribute to the individuation of the nature of perceptual experiences – between perceptual experiences and features of a subject's environment, even if these are reliable, cannot contribute to the subject's having justified beliefs formed on the basis of these experiences.

The question, of course, is why this should be so? Unfortunately, at least as far as I can tell, Burge does not provide any further defence of this assumption. However, rather than questioning it, I want to proceed by granting Burge his key thought about when the reliability of perceptual experiences contributes to the

epistemic efficacy of these experiences and when it does not. For a further claim that Burge makes, which I will eventually question, is not only that reliable connections between perceptual experiences and features of the world fail to contribute to the epistemic efficacy of these experiences if these connections are merely accidental relative to the nature of the experiences, but also that the unreliability of perceptual experiences in a given environment need not affect their epistemic efficacy if the environment is not the one relative to which the nature of these perceptual experiences is individuated. Thus:

Perceptual errors or unreliabilities that are perceptually indiscernible and derive from brute abnormalities in the environmental conditions that cause perceptual states do not undermine warrant. Both unreliability and reliability in conditions other than those that played a role in explaining the nature of the perceptual state and the exercise of the perceptual competence are accidental relative to those natures. So reliability and unreliability in such conditions are irrelevant to the connection between warrant and veridicality. The only reliability that is relevant to the contribution of perceptual states to perceptual warrant is one that attaches to the states' normal functioning in the conditions that explain their natures. That is the only reliability that is non-accidental relative to the natures of the perceptual states. (ibid, 536)

So in looking at Burge's view, we are now confronted with two different claims. To repeat, the first claim is that if perceptual experiences provide one with epistemic justification for believing propositions about one's surroundings then the nature of these experiences has to be connected in a reliable way to features of the

environment relative to which the nature of these experiences is individuated. The second claim is that perceptual experiences can provide one with epistemic justification for believing propositions about one's environment even if these experiences are highly unreliable in that environment provided that the environment in which the experiences are unreliable is not the one relative to which the nature of the experiences is individuated.

One question is how these two claims are related and, in particular, whether or not they are independent from one another. Burge thinks that they are not. This is made clear in the following passage:

... the normal environment – the conditions in which content is explained and established – is privileged in explanation of entitlement. Its privilege derives from the fact that it plays a central role in making the individual's states what they are – a role that abnormal environments do not play. This privilege entails the irrelevance to entitlement of reliability and unreliability in other conditions. So it extends to indiscernible abnormal environments that the individual might contingently find himself in. (ibid, 536)

The line of argument presented here can be summarised as follows. (A) The normal environment is privileged with respect to the explanation of perceptual justification (or entitlement); it is privileged in the sense that only reliable connections between features of the normal environment and a subject's perceptual experiences can contribute to epistemic efficacy of these perceptual experiences (this is essentially the first claim above). From (A) it follows that (B) reliable connections between features of the world and a subject's perceptual experiences outside of the subject's

normal environment contribute nothing to the epistemic efficacy of these experiences. It also follows from (A) that (C) the unreliability of perceptual experiences in a given environment other than the normal environment does not affect the epistemic efficacy of these experiences (provided of course that these experiences are reliable in the subject's normal environment)

Now I take it that the entailment from (A) to (B) is unproblematic. But why should we accept that (A) entails (C)? Unfortunately again, Burge does not seem to offer a defence of this crucial move. At the very least, it is not immediately clear why (A) should entail (C). For we could accept the proposition that perceptual experiences provide justification only if they are reliable relative to the subject's normal environment and yet insist that they provide one with justification only if one is actually in environments in which the experiences are reliable. Call this conjunction proposition (D). One of the reasons one might hold (D) to be true is that it is only in environments in which perceptual experiences are actually reliable that they are, objectively speaking, a good route to truth. (D), in virtue of its second conjunct, is of course inconsistent with (C). So if, as Burge says, (A) entails (C), the second conjunct of (D) must be false since (A) entails the first conjunct of (D). But *prima facie*, there does not seem to be any inconsistency in (D).

Leave this issue aside for the moment (I will return to it below). Of course, if we were to accept (C) (that is, the second of the two claims of Burge's listed earlier), then we can give (at least a partial) explanation of the new evil demon intuition. A victim of an evil demon can have justified perceptual beliefs insofar as he forms these beliefs on the basis of perceptual experiences with representational contents that are reliably veridical in his normal environment even if these contents are highly unreliable in the environment he inhabits. On Burge's account, evil demon victims

are entitled to rely on their perceptual experiences in forming beliefs about their surroundings insofar as these experiences are by their nature reliably connected to the truth of these beliefs, not in his current environment but in the environment in which his perceptual system has developed.

Evaluating Bergman's and Burge's Views

Bergman and Burge each have now presented us with two different objectivist views of how perceptual beliefs can be epistemically justified for us. In evaluating their views, I want to focus specifically on the question of whether or not they successfully accommodate the new evil demon intuition.

I will begin with Bergman's proper function account. To recall, this view holds that the victim of an evil demon is epistemically justified in holding his perceptual beliefs, even if these are reliably false, because these beliefs are the result of a cognitive process that is functioning properly, designed to produce true beliefs, and reliable in the environment for which it is designed. But why should we think that the conclusion that the evil demon victim has justified perceptual beliefs should follow from these conditions?

On brief reflection, this conclusion does not seem to follow naturally at all. Consider, for example, this analogy. Suppose that I have the intention of taking part in a bicycle race in a mountainous and rocky terrain. The bike that I choose for doing so, however, is a highly specialised and sophisticated racing bike that was used by eight-time Tour de France winner Lance Armstrong on the last Tour. This bike is in excellent working order and as such functioning as it is supposed to function, is

designed to maximise a rider's ability to ride at very high speeds on paved and reasonably smooth roads and does so reliably under the conditions for which it was designed. Of course, given the thinness of the wheels on this bike, I struggle to make my way through the race at a high speed without damaging the wheels in the rocky terrain that I am riding in. Given the environment that I am in, it is clear that it is entirely inappropriate for me to be using this bike. Instead, I should have opted for a specialised mountain-bike with the kind of wheels, construction and design appropriate for racing in a mountainous and rocky terrain. In evaluating the appropriateness of using the racing bike, it is simply entirely irrelevant whether or not the bike is functioning as it is supposed to function, whether or not it is designed to maximise the rider's ability to achieve high speeds on normal roads, and whether or not it reliably does so under the conditions for which it was designed.

By analogy, we might thus say that it is epistemically entirely inappropriate, objectively speaking, for the victim of an evil demon to rely on belief-forming processes that are highly unreliable in the environment that he finds himself in. And the appeal to the fact that these belief-forming processes are functioning properly, are designed to produce true beliefs and are reliable in the sort of environment for which they were designed seems to be entirely irrelevant.

Of course, in the case of evaluating the epistemic appropriateness of relying on a particular belief-forming process, we have to recall that such evaluations depend in a crucial way on what epistemic goal is meant to be satisfied. Thus there may well be good reasons for thinking that we shouldn't simply evaluate the epistemic appropriateness of a subject's relying on a belief-forming process in terms of whether it reliably leads to true beliefs in his particular environment but rather whether it reliably leads to true beliefs in a wider reference class of environments.

One good reason for evaluating a subject's reliance on a belief-forming process in this way is if we can suppose that this subject actually conducts (or perhaps, is physically capable of conducting) his epistemic affairs in a wider reference class of environments.

However, this is not a strategy that Bergman can fruitfully pursue. For it is entirely possible for there to be a victim of an evil demon who conducts his epistemic affairs entirely in an environment or in a class of environments in which his perceptual belief-forming processes are highly unreliable but, still, are functioning properly, designed to produce true beliefs and reliable in the environment for which it was designed. Thus suppose that the evil demon creates a world much like our world. As inhabitants of this world, he creates organisms with the same cognitive processes that we have. In particular, he designs these organisms in such a way that they have perceptual belief-forming processes that are aimed at truth and are in fact (globally) reliable in this world. The organisms heavily rely on these processes in forming beliefs about their surroundings and we can assume, here, that it is entirely appropriate, epistemically and objectively, for them to do so given that these processes produce mostly true beliefs. However, one day the demon cruelly destroys this world and replaces it with a world radically different from what it was but continues to feed the organisms he created with the same kind of perceptual experiences they were enjoying before. The perceptual belief-forming processes of these organisms are now (globally) highly unreliable. But should we say that the organisms' perceptual belief-forming processes continue to produce justified beliefs for these organisms?

From an objective perspective, the answer to this question seems to be 'no'. After all, the reliance on their perceptual belief-forming processes does not leave

these organisms with beliefs that are mostly true, even in a global context. Thus relative to the goal of forming mostly true beliefs in conducting their epistemic affairs, reliance on their perceptual belief-forming process is, objectively speaking, entirely inappropriate. Again, the fact that these processes are functioning properly, designed to produce true beliefs and reliable in the environment for which they are designed seems entirely irrelevant. For there does not seem to be any relevant epistemic goal here relative to which the reliance on these processes can be, objectively speaking, epistemically appropriate.

Burge's view ultimately faces problems that are similar to the ones faced by Bergman's view. Recall that, according to Burge, a subject is epistemically justified in relying on his perceptual experiences insofar as these experiences are by their very nature connected with reliable representation. This means that a subject is epistemically justified insofar as these perceptual experiences have the nature that they do as a result of developing in an environment in which these experiences are reliably veridical. Furthermore, Burge holds that it is only insofar as the reliability of perceptual experiences is constitutively connected with the nature of these experiences that it contributes to the justification of beliefs based on them. Crucially, this last thought is supposed to entail that reliance on perceptual experiences that are by their very nature connected to reliable representation can be justified for a subject even in environments in which these experiences are highly unreliable.

As we have seen, Burge does not provide any further explanation for why this entailment is supposed to hold and earlier I raised the worry that the supposed entailment was far from obvious. We have now also seen, in relation to Bergman's account, that the notion that a process or ground is by its nature connected in some way to reliably true beliefs is not sufficient for such a process or ground to provide a

subject with epistemic justification. After all, Bergman's properly functioning perceptual processes are by their nature reliably connected to truth – that is they are aimed at truth and reliable in the environment for which they are designed – but this was not enough to explain how they could produce justified beliefs in relevant evil demon environments.

Returning to Burge's view, suppose we have a victim of an evil demon who has perceptual experiences that are by their nature connected with reliable representation in their home environment but that are highly unreliable in the environment the victim is now inhabiting. It is clear that in one sense it is epistemically inappropriate, objectively speaking, for this subject to rely on his perceptual experiences given that they are largely unveridical in his environment and therefore tend to produce false beliefs. At this point, the epistemic relevance question applies once more: why should it be epistemically appropriate for a subject to rely on perceptual experiences that are reliable in the environment in which their nature has developed but highly unreliable in the environment he currently inhabits? The reliability of these experiences in their home environment seems to be of no relevance at all.

Of course, once again, we have to note that evaluations of the epistemic appropriateness of relying on particular processes or grounds for belief are highly dependent on the goal relative to which we should make such evaluations. Thus we might say, as before, that the objective appropriateness of relying on specific grounds or belief-forming processes should be judged on the basis of the way a subject actually conducts his epistemic affairs. And if a subject conducts his epistemic affairs within wider class of different local environments rather than within a local environment, then the objective appropriateness of relying on specific

grounds or belief-forming processes would have to be judged in terms of whether these grounds or processes are reliable relative to the wider class.

However, as with Bergman, this point is of no use to Burge's view. Thus suppose a demon ensures that a subject's perceptual system develops in a global environment in which the subject's perceptual experiences are reliably veridical. This subject's perceptual states will have the contents that they do in virtue of these globally reliable connections. Suppose, finally, that the subject does in fact conduct his epistemic affairs within the global environment in which his perceptual experiences have developed. Cruelly, the demon now destroys this global environment, replaces it with a very different global environment but continues to feed the subject with the same kind of perceptual experiences that he was enjoying before. These perceptual experiences are now highly unreliable within the subject's global environment. The question now is whether it is objectively appropriate for the subject to rely on his perceptual experiences, given the aim of generally having true beliefs and avoiding false ones. With respect to the given cognitive aim, it seems clear that reliance on perceptual experiences is objectively inappropriate – that is, such reliance does not objectively satisfy the given cognitive aim. This is so even if these perceptual experiences are by their nature connected to reliable representation in their home environment.

The problem with both Bergman's and Burge's views is thus essentially the same. As we have seen, both these theorists attempt to preserve an objectivist account of perceptual justification while at the same time trying to accommodate the new evil demon intuition by insisting that perceptual grounds or processes need to be reliable only relative to a specific environment but not relative to the environment the subject happens to inhabit. But neither of them makes clear why it should be

deemed epistemically appropriate, from an objective point of view, to rely on perceptual grounds or processes that are reliable only relative to the specific environments they suggest. As we have seen, whether reliance on specific grounds or processes is, objectively speaking, epistemically appropriate is a matter of whether or not such reliance satisfies a specific epistemic goal. But we have also seen that there can be cases in which it is not at all clear what epistemic goal is supposed to be satisfied by relying on a belief-forming process that is functioning properly, aimed at truth, and reliable in the environment for which it was designed or by relying on perceptual experiences that are by their nature reliably connected to veridical representation. That is why the views that Bergman and Burge develop ultimately are highly problematic.

Chapter 2

Introduction

In the last chapter, we investigated different proposals for how the Basic Principle about Perceptual Justification could be explained in objectivist terms and found no objectivist account that could adequately accommodate the intuition that the victim of an evil demon, whose perceptual experiences do not reliably produce true beliefs about his environment but are otherwise subjectively indistinguishable from the experiences that we enjoy, could nonetheless have perceptual beliefs about his environment that are epistemically justified for him. If we cannot in fact adequately explain the Basic Principle in objectivist terms, the natural move to make is to try to determine instead whether we can have more success in seeking an explanation of the Basic Principle in terms of the subjective conception of epistemic justification. I will refer to accounts that seek to explain the Basic Principle in terms of the subjective conception of epistemic justification as *subjectivist* accounts.

On one such a conception, a subject S has epistemic justification for believing p only if S has some appropriate belief to the effect that conditions obtain that make it more likely than not that p is true. Call this the *subjective-doxastic* conception of epistemic justification. Crucially, having an appropriate belief to the effect that conditions obtain that make it more likely than not that p is true need not entail that those conditions actually obtain, and this is a feature of the subjective conception of

epistemic justification that, given the new evil demon intuition, makes it a prima facie attractive option in terms of which the Basic Principle can be explained.

Alternatively, we might also think of a perceptual experience itself as constituting a subjective perspective on how things are in the world. Call this the *subjective-experiential* conception of epistemic justification. Thus, it is because the subject, in having the perceptual experience, occupies a subjective perspective from which conditions obtain that make it more likely than not that a proposition *p* is true that he is justified in forming the belief that *p* on the basis of his having the experience. In fact this is roughly the explanation of the Basic Principle that I will be proposing in this chapter. Before doing so, however, I want to first turn to the question of whether we can explain the Basic Principle in terms of the subjective-doxastic conception of epistemic justification.

In order to answer this question, we need to turn briefly to the notion of an ‘appropriate belief’ that figures in our formulation of what the subjective-doxastic conception of epistemic justification amounts to. What makes a belief to the effect that conditions obtain that make it more likely than not that a proposition *p* is true an appropriate one such that this belief can constitute a subject’s justification for believing that *p* is the case? Obviously, it is not sufficient for a subject to have a justified belief that *p* that he merely has a belief to the effect that the relevant conditions obtain. Suppose, for example, that *S* forms the latter belief on the basis of reading tea leaves. We can assume that unless *S* has some good reasons for thinking that reading tea leaves reliably leads to true beliefs, *S* need not have any justification for believing that *p* in this case. The question is therefore what further properties a belief needs to have in order for this belief to be appropriate in the relevant sense for constituting a subject’s justification for another belief.

Call an appropriate belief of this kind a *justifying belief*. In specifying the required further properties that a justifying belief needs to have, we might insist that justifying beliefs themselves need to be epistemically justified for the subject in some way. However, it is doubtful that this would lead us to a viable subjectivist explanation of the Basic Principle. To see this we can assume that if the Basic Principle is explained on the basis of S's having some belief to the effect that a situation obtains which makes it more likely than not that p is true this justifying belief is a belief about the experience he is undergoing. Suppose then that we say that a subject S's belief that an x is F, in circumstances in which he has a perceptual experience of an x being F, is justified for S in virtue of his believing that he is having an experience of an x being F and that having such an experience makes it more likely than not that an x is F. In what way can S's justifying beliefs themselves be justified in this case? We can assume for our purposes here that S's first justifying belief – that he is having an experience of an x being F – can be justified for him in a relatively straightforward way. For example, we might say that this belief is objectively justified on the basis of his actually having an experience of the sort he believes to be having and thereby justified on the basis of a ground that entails the truth of his belief. The difficulty, however, lies in trying to give an account of how S's second justifying belief can be justified for S. The issues here are familiar. For one, the proposition that he is having an experience of an x being F does not entail that an x is actually F. Hence, we cannot explain S's justification for believing that having an experience of an x being F makes it likely (or rather, entails) that an x is F in terms of, say, S's grasp of the concept of an experience of an x being F. So how else can the second justifying belief be justified for S? A plausible assumption is that, if this belief is justified at all for S, it is justified by S's having further justifying

beliefs. Thus one option is to say that S's belief that his having an experience of an x being F makes it more likely than not that an x is F is justified on the basis of his justifiably believing that past instances of his having an experience of this kind were more often than not correlated with the fact that an x is F. But it is hard to see how S's belief about the relevant past instances could themselves be justified for S unless he already had justified beliefs to the effect that an x is F in circumstances in which S had perceptual experiences of an x being F. And, of course, these are the kind of beliefs whose justification we are trying to account for. Another option is to say that S's belief that his having an experience of an x being F makes it more likely than not that an x is F is justified on the basis of his justifiably believing that the appeal to an x actually being F is the best explanation for why S is currently having an experience of an x being F. However, it has proved notoriously difficult to explain how such inferences to the best explanation can be justified for us. Furthermore, it is far from clear, even if it is true that an x actually being F is the best explanation for why a subject is having an experience of an x being F, why this should entail that one's having this experience makes it more likely than not that an x is actually F.

Given these difficulties, what other options are left? One further alternative is to say that a subject S's belief, in circumstances in which he is having a perceptual experience of an x being F, is justified for S in virtue of his believing that he can see that an x is F and that seeing that an x is F entails that an x is F. Thus in this case, S's belief that an x is F is epistemically supported by his having a belief that he is having a factive perceptual experience which entails that an x is F. Here we can assume that S's belief that seeing that an x is F entails that an x is F can be straightforwardly justified for S. For example, we might simply say that this belief is justified for S on the basis of S's grasp of the concept of the state of seeing that an x is F. The

difficulty with this alternative, however, lies in trying to give an account of how the S's first justifying belief – that he sees that an x is F – can be justified for S. Thus, on the one hand, we might try to explain S's justification for this belief in objectivist terms. For example, we might say that S's belief that he sees that an x is F is justified for S when he bases this belief on his having an experience of an x being F which is a reliably correlated with his seeing that an x is F. But objectivist accounts like this could not explain how victims of an evil demon might be justified in having such a belief. On the other hand, we might try to explain S's justification for believing that he sees that an x is F in subjectivist-doxastic terms. But then the question is what justifying beliefs S could have for believing that he sees that an x is F. Again, it is difficult to see how S could have justifying beliefs for believing that he sees that an x is F without his already having some justified belief that an x is F.

It is difficult to see, therefore, how the Basic Principle could be explained in subjectivist-doxastic terms on the assumption that justifying beliefs themselves need to be justified. Instead of investigating this issue any further, I want to explore in more detail a way developing a subjectivist-doxastic account of the Basic Principle that does not conceive of justifying beliefs as themselves needing to be justified. One such account can be derived from Richard Foley's conception of an epistemically rational belief and this is the account with which I will begin. However, as I will be showing, this conception faces some serious difficulties. I will then turn to the view on perceptual justification that is offered by John McDowell, which I will suggest can be understood as presupposing, in part, a subjective-experiential conception of epistemic justification. This view, too, suffers from significant defects but the discussion of McDowell's view helps to illuminate a conception of perceptual justification that can underpin a viable explanation of the Basic Principle. Crucially,

this explanation of the Basic Principle will appeal specifically to the phenomenal character of perceptual experiences.

Foley's Conception of an Epistemically Rational Belief

In his book *The Theory of Epistemic Rationality*, Foley develops a theory of what constitutes an epistemically rational belief. In what follows, I will try determine whether this notion of an epistemically rational belief can provide us with an understanding of how our perceptual beliefs can be justified in the way that the Basic Principle states that they can be.

Foley's theory is a version of an Aristotelian view of rationality, which in general conceives of rationality in terms of a subject's goals and what this subject would consider to be an effective way to pursue these goals upon careful reflection. Epistemic rationality as a particular form of rationality in the Aristotelian sense is then defined in terms of a specifically epistemic goal, which as Foley conceives of it, is the goal of now having true beliefs and avoiding false ones. Given this specifically epistemic goal, Foley suggests the following account of what constitutes an epistemically rational belief:

...it is epistemically rational for an individual to be persuaded of the truth of just those propositions that are the conclusions of arguments that he would regard as likely to be truth-preserving were he to be reflective and that in addition have premises that he would uncover no good reason to be

suspicious of were he to be reflective. (Foley, *The Theory of Epistemic Rationality*, 1987, p. 5)

The notion of an epistemically rational belief is thus defined in terms of a certain counterfactual claim. On this view, a belief that *p* is epistemically rational for a subject *S* if there is an argument that *S* *would* regard as establishing *p* and whose premises he *would* accept if he *were* to be reflective.

One question that this account immediately raises is what it is for a subject to be reflective in the relevant sense. In one relevant sense, for a subject to be reflective at a time *t* is for him to bring bear on the questions of whether or not an argument *a* is truth-preserving and of whether or not the premises in *a* are such that there are no good reasons to be suspicious of them, other relevant propositions that *S* believes at *t*. Thus *S* might consider whether there are further arguments which he believes to be truth-conducive and whose premises he would accept upon further reflection which either support or undermine the belief that *a* is truth-conducive and/or the belief that the premises in *a* are true. Ultimately to Foley, however, reflecting on the questions of whether or not an argument *a* is truth-preserving and on whether or not the premises in *a* are such that they are no good reasons not to accept them will bring to bear a subject's deepest epistemic standards. These, in turn, are a 'function of what he now is disposed to think about [an argument] were he to be sufficiently reflective.' (ibid, p. 35) Thus at some point in one's reflection, one will simply be disposed to take an argument *a* as truth-conducive without there having to be yet a further argument which, on yet further reflection, one would take to support the assumption that *a* is truth-conducive. Here a subject comes to a point at which his own deepest epistemic standards are simply made manifest: '...if an individual on

reflection is disposed to think that the argument is sufficiently likely to be truth preserving and if in addition he is so disposed that further reflection would not change his mind about this, then the argument conforms to his own deepest epistemic standards.’ (ibid, p. 35) Epistemic rationality, on the Foleyian view, is thus entirely a matter of a subject’s reflective perspective. That is, it is matter entirely of whether or not a subject’s epistemic goal of believing p only if p is true is met from the subject’s own reflective point of view. Crucially, being epistemically rational in this sense does not entail that a subject systematically achieves his goal of now believing truths and avoiding falsehoods. In fact, it is entirely consistent with being epistemically rational in the Foleyian sense that a subject has epistemically rational beliefs that are mostly false.

It is also consistent with the Foleyian view that a subject can be epistemically rational even if, on reflection, he would be persuaded of the truths of propositions on the basis of arguments in a way that would seem to most ordinary people to be utterly bizarre. This is because subjects may ultimately have radically different deep epistemic standards. Thus some subjects might on deep reflection be disposed to judge arguments as being truth-preserving and accept the premises of these arguments as true in a way that others might find to be radically misguided. (I will return to this issue below) Finally, I take it that being epistemically rational in the Foleyian sense does not entail that a subject’s own reflective perspective from which it will seem to him as if his epistemic goals are sufficiently met is itself epistemically rational for him. Thus, a subject can, upon sufficient reflection, be persuaded of the truth of a proposition p by arguing to p from premise q without there being a further argument, which upon yet further reflection, the subject would take to support q and to support the inference from q to p . As we have seen, his acceptance of q and of the

validity of the inference from q to p can simply be a manifestation of his own deep epistemic standards. Furthermore, his acceptance of q and of the validity of the inference from q to p need not be epistemically justified for S in any other sense. For example, it need not be the case that S's reflecting on q and on the validity of the inference from q to p is such that it reliably produces true beliefs in regards to q and the inference from q to p.

Given this last point, we might of course ask why we should think that the kind of subjective reflective perspective that Foley appeals to should generate any form of epistemically rational or justified beliefs. In other words, why should the epistemic rationality of believing a proposition p be a matter of certain kinds of further beliefs that a subject would adopt if he were to engage in a certain kind of reflection, if these beliefs themselves need not be epistemically rational for S or justified for him in any other way? This of course relates to the general question raised earlier of what properties make a belief a justifying belief. Foley's answer to this question is that the epistemic rationality of believing a proposition p is a matter of certain beliefs that one would adopt on the basis of reflection because such beliefs render the belief that p 'invulnerable to intellectual self-criticism':

Being rational in this sense involves making oneself invulnerable to intellectual self-criticism to the extent possible. It is a matter of having opinions and using faculties, methods, and practices capable of standing up to one's own, most severe scrutiny. For an opinion to pass this test, it must not be the case that one's other opinions could be used to mount what one on reflection would regard as a convincing critique of it. Nor can it be the case that one has, or would have on reflection, criticisms about the ways one has

acquired the opinion. (Foley, *Epistemic Rationality as Invulnerability to Self-Criticism*, 2003, p. 3) .

Thus, according to Foley, subjecting one's beliefs to the kind of reflection he outlines constitutes a most severe kind of intellectual self-scrutiny of those beliefs. Beliefs that are subject to, and survive, this kind of scrutiny count for him as being epistemically rational.

Now, I think we can grant Foley the point that, at the very least, his conception does pick out a certain form of epistemic rationality that beliefs can have. The question of interest to us, however, is whether this conception provides us with a means of understanding how our perceptual beliefs can be epistemically justified for us. So how can this conception of epistemic rationality provide us with an explanation of the Basic Principle about Perceptual Justification? In the first instance, a Foleyian explanation of the Basic Principle will say that that a subject, when having a perceptual experience of something's being F, will be justified in Foley's sense of being epistemically rational in believing that something is F insofar as there is an argument that has the proposition that something is F as its conclusion, which S would take to be truth-preserving and whose premises S would accept were he to be sufficiently reflective. However, given Foley's conception of epistemic rationality, there are in principle many different ways in which subjects, when they are having a perceptual experience of something's being F, can be epistemically rational in believing that something is F. For example, one subject might be epistemically rational in believing that something is F in circumstances in which he is having the relevant sort of perceptual experience if, on reflection, he would be persuaded of the truth of the proposition that something is F, on the basis of arguing

from the premise that he is having an experience of something's being F to the conclusion that something is F. Another subject, however, might be epistemically rational if, on reflection, he would argue to the conclusion that something is F from the premise that he can see that something is F. And so on. That is because, as we have seen, deep epistemic standards can vary among different subjects.

What should we make of Foley's conception of epistemic rationality and its application to the Basic Principle? To answer this question, it will be helpful to consider some prominent objections to Foley's general view. One line of criticism is that Foley's conception of epistemic rationality does not really pick out an interesting or important epistemic property. Richard Feldman expresses this worry in the following passage:

I don't see why one's own deepest epistemic standards identify any particular interesting or important concept of rationality. This point is best raised as a rhetorical question: Why think my deep standards are so important for evaluating my rationality? Why not use your standards, or my Uncle Al's, standards, or community standards,...? Foley seems to acknowledge that there is a sense of 'rational' for each such standard. I'm suspicious of the idea that there are so many senses of the term, but if there are so many senses, I'm especially suspicious of the idea that the one he singles out is of any great importance or that it somehow identifies a central epistemic concept. There is, I admit, something virtuous about conforming to one's own standards. It's hypocritical, or dishonest, not to. But I'm not sure that the relevant virtue is appropriately characterized as being rational. (Feldman, 1989, p. 157)

One question that Feldman's criticism here raises, however, is what exactly makes an epistemic concept an interesting or important one. Perhaps, we can view the matter in this way. Suppose we begin with the idea that we have some pre-theoretical understanding of the kinds of beliefs that are epistemically justified for us. Thus, pre-theoretically we think that our perceptual beliefs are generally epistemically justified for us and the same applies to beliefs that are based on our memories or the testimony of others. Given this, we might say that an interesting or important epistemic notion is one that can be applied to these beliefs and thereby illuminate in what sense they are epistemically justified for us. If this is the sense in which an epistemic notion is interesting or important, I do think that the charge that Foley's notion of epistemic rationality is ultimately uninteresting or unimportant, at least insofar as our perceptual beliefs are concerned, is basically sound, but we need to take some care in spelling out the argument in support of this contention.

Now one way of not arguing for this contention, however, is by presupposing, implicitly or explicitly, that an interesting or important epistemic notion can only be one that renders our beliefs justified in an objective sense. Thus take the following criticism of Foley's conception of epistemic rationality by William Alston:

The key to the whole matter is the fact that, as Foley acknowledges and even insists, the basic epistemic goal is 'now to have true beliefs and now not to have false beliefs'. Hence our interest in the concept of the rationality of belief, or more generally in a mode of epistemically evaluating beliefs, will be in direct proportion to the relevance of this goal. Most directly relevant is the evaluation of beliefs as true or false. Next down the line comes the consideration of the likelihood of a belief's being true, relative to some

publicly accessible body of evidence, or relative to what the subject has or to what is available to him, or relative to the basis of the belief...At the bottom of the list comes a purely subjective conception, whether S believes that it is acceptable for him to believe that p. And just above that, but significantly below all the other alternatives listed, is F-rationality, whether S would take it, on reflection, that he has an uncontroversial argument for p. This is of little interest, relative to the epistemic goal, just because it has no determinate implication, in itself, for (the likelihood of) the truth or falsity of the belief. (Alston, 1989, pp. 146-7)

According to Alston, relative to the epistemic goal of having true beliefs and avoiding false ones, the objective conceptions of epistemic justification or rationality are much more important and interesting than the subjective conception that Foley develops. The problem with the latter, so Alston argues, is precisely that it fails to ensure that our justified or rational beliefs are more likely than not to be true. The problem with this objection to Foley's view, however, especially when the issue turns to our perceptual beliefs, is that it may simply not be possible to account for some of our justified beliefs in objectivist terms. In the above passage, Alston suggests that there is a certain hierarchy of different ways of epistemically evaluating a belief which represents the relative importance of a particular conception of epistemic justification or rationality. At the top of this hierarchy we find various objectivist forms of epistemic evaluation, while the subjectivist ones are further down the list. However, Alston does not mention what forms of epistemic evaluations might fall between the objectivist ones he lists and Foley's subjectivist conception of epistemic rationality. So for all we have been given, and on the

assumption that we cannot explain how perceptual beliefs in particular are justified for us in objectivist terms, Foley's conception of epistemic rationality may well be the best form of epistemic evaluation we have to explain the Basic Principle.

A more promising approach to showing that Foley's conception of epistemic rationality is uninteresting or unimportant in the relevant sense is taken by Marshall Swain in his criticisms of Foley's views. One point on which these criticisms focus is, as we have seen, that Foley allows for the possibility of cases when a belief is epistemically rational for a subject, even if, on reflection, he would argue for the belief on the basis of what we would consider to be very bizarre. Swain provides the following example

When confronted with the premise $p \vee (p \ \& \ q)$...an individual might correctly conclude p , but only because one his deepest standards directs him to derive p from any disjunction with p as a single disjunct. If confronted with $p \vee q$, he would also conclude p ...This person is just not very good at logic, and thinks that lots of inferences are valid when they are not. In what sense are such a person's beliefs epistemically rational? (Swain, 1989, p. 164)

In response to Swain's question at the end of this passage, and as I suggested earlier, I think that we can grant that there is a perfectly good sense in which the person's belief is epistemically rational, which is simply the sense of being epistemically rational under Foley's conception. However, even if we do grant this, we certainly still want to say that a subject who can logically reason in a competent way is epistemically rational in a very different and, we might say, more interesting and important sense. The intuition concerning the subject in Swain's example and the

subject who can logically reason in a competent way is thus that the latter subject is, epistemically speaking, much better off than the former. Swain himself seeks to distinguish these two subjects by insisting, like Alston, that genuine epistemic rationality – or in our terms, a more interesting or important sense of it – requires an objectivist connection to truth. That is, to Swain, if epistemic rationality is to be understood in terms of adopting only those beliefs that would survive deep reflection in Foley’s sense, then it should follow that those beliefs are more likely than not to be true. But once again, we have to note that the insistence on an objective truth-connection in this way is unlikely, given the new evil demon intuition, to yield a viable conception of how our perceptual beliefs can be epistemically justified for us. Nonetheless, I think that there is an important point that has emerged from Swain’s example and this is that our perceptual beliefs can be epistemically justified or rational for us in a way that makes them much better, epistemically speaking, than some beliefs that are rendered perfectly rational by Foley’s conception of epistemic rationality. Thus, like the subject who can logically reason in a competent way, I suggest that a subject who forms a belief about his surroundings on the basis of the relevant sort of perceptual experience can be justified in having this belief in a sense that is much more interesting or important than the sense in which the subject in Swain’s example is epistemically rational. The challenge, therefore, is to develop a notion of epistemic justification or rationality that can make the right distinctions between the different subjects we have considered and it seems clear that Foley’s subjectivist conception is incapable of making those distinctions.

McDowell's Conception of Perceptual Justification

As I suggested earlier, we may be able to provide an explanation of the Basic Principle by drawing on certain elements that can be found in McDowell's conception of perceptual justification. Before, turning to this explanation, however, we first need to consider McDowell's conception in its entirety.

Some of the central elements of McDowell's conception of perceptual justification are outlined in the following passage from his *Mind and World*:

In a particular experience in which one is not misled, what one takes in is *that things are thus and so*. *That things are thus and so* is the content of the experience, and it can also be the content of a judgment: it becomes the content of a judgment if the subject decides to take the experience at face value. So it is conceptual content. But *that things are thus and so* is also, if one is not misled, an aspect of the layout of the world: it is how things are. Thus the idea of conceptually structured operations of receptivity puts us in a position to speak of experience as openness to the layout of reality. Experience enables the layout of reality itself to exert a rational influence on what a subject thinks. (McDowell, *Mind and World*, 1994, p. 26)

There are in fact several elements here that make up McDowell's conception of perceptual justification. However, the key element to McDowell's thought seems to be that perceptual experiences can in some cases be 'an openness to the layout of reality', that is, to mind-independent, environmental facts. What follows from this, or so McDowell argues, is that in those cases in which perceptual experiences do

amount to an openness to facts, these facts themselves can ‘exert a rational influence’ on our thoughts. The idea that McDowell seems to be expressing here is that, given the right kind of perceptual experience, an environmental fact that *p* can be one’s *reason* for believing the proposition that *p*. The right kind of perceptual experiences can therefore provide us with justification for believing propositions about the external world by providing us with, or bringing into view, the facts that make these propositions true. In providing us with such facts, or in bringing them into view, these experiences provide us with reasons for beliefs about our environment.

Stated thus far, McDowell’s conception of perceptual justification suggests the following explanation of the Basic Principle: when having a perceptual experience of something’s being *F* and forming the belief that something is *F* on the basis of having this experience, this belief is *prima facie* justified for a subject to the extent that the given perceptual experience amounts to an openness to the fact that something is *F* and, therefore, to the extent this experience brings into view a mind-independent fact that acts as the subject’s reason for forming the belief that something is *F*. Now the important first point to note about this explanation of the Basic Principle is that it is of an objectivist sort: that is, the Basic Principle is explained on the basis of the thought that a belief that *p* is perceptually justified for a subject insofar as he bases this belief on his having a perceptual experience that objectively entails, or is necessarily correlated with the truth of, the proposition that *p*.

However, it is equally important to note that this objectivist element in McDowell’s picture of perceptual justification, if it plays a role at all, forms only part of the picture as a whole. In its complete form, McDowell’s view on perceptual

justification crucially also involves a subjectivist element. His insistence on this subjectivist element derives from his acceptance that an epistemically justified belief (and ultimately knowledge) essentially belongs to *the space of reasons*. Thus to him, one has an epistemically justified belief that p only if one has a reason for believing that p is true. What we need to determine here is how exactly McDowell understands the notion that when one has a perceptually justified belief that p one has a perceptual reason for believing that p is true.

On one general understanding of the notion of an epistemic reason, one has a reason for believing a proposition p when one can cite a valid argument consisting of premises q, r, s, and so forth that has p as its conclusion. Of course, further conditions would have to be met for one to have genuine justification for believing p on the basis of one's ability to cite a valid argument in support of believing p in this way. For example, we might suppose that these conditions include that one has justification for believing each of the premises as well as justification for believing that the premises validly entail the conclusion. Having reasons in this sense can often play an important dialectical role. Thus one can appropriately cite reasons of this kind in justifying one's beliefs to others or in trying to convince them to adopt the same belief in the conclusion. Again, there will be further conditions on when citing an argument in support of a belief in this way can be dialectically effective.

A reason for believing p on the foregoing understanding is thus essentially a proposition (or a fact) that one can cite in support of one's believing that p. It is what Hannah Ginsborg, in her recent discussion of McDowell's view of perceptual justification, calls a reason¹ for believing a proposition p. As Ginsborg puts it, this sense of 'reason' "is the one we most naturally invoke when we are concerned with the first-person perspective from which a subject assesses her beliefs or potential

beliefs and considers which ones she ought to retain or adopt.” (Ginsborg, 2006, p. 290) It is important, however, to carefully distinguish between a reason qua proposition or fact in this sense on the one hand and the mental state in virtue of which a subject possesses such a reason on the other. Thus, as it was explained in the previous paragraph, for one to have a reason *r* qua proposition for believing that *p*, it is intuitive to suppose that one needs to have some justification for believing *r* as well as for believing that *r* entails *p*. This would mean in turn that one has a justified belief that *p* on the basis of one’s having a reason *r* qua proposition for believing *p* only if one believes that *r* and has some justification for doing so.

Now, crucially, we might consider the justified belief that *r* which, in this case, constitutes a subject’s justification for believing *p* to be the subject’s ‘reason’ for doing so. However, the sense in which a mental state such as a belief can be a subject’s reason for believing another proposition is different from the sense in which a proposition can be a reason for him. To mark the difference between these two senses, Ginsborg introduces the notion of a reason² for believing something and defines this notion as follows:

The second sense of reason...is the sense we invoke when, from a third-person perspective, we assess the rationality of someone else’s beliefs, or, relatedly, try to make her beliefs rationally intelligible. From this perspective, in contrast to the first-person perspective occupied by the subject, the actual facts are irrelevant to the determination of what reason the subject has and whether they are good reasons. As a subject assessing my own beliefs, what I need to determine is what the facts are independently of those beliefs: if the issue is whether I am justified in believing that it has rained, I need to

determine whether the streets are wet, not whether I believe that the streets are wet. But if I am assessing someone else's beliefs, then I need to determine how things present themselves as being from her point of view....[we] specify a psychological state, typically another belief, in the light of which her original belief can be recognised from a third-person perspective, as rational. (Ibid, p. 290)

Thus a mental state M of a subject S can be a reason² for having a belief B when, from the third-person perspective, we can make S's having B rationally intelligible in light of his having M. Crucially, Ginsborg thinks that the notions of a reason¹ and a reason² are conceptually related. This means, in particular, that a subject can have a reason² for believing p only he takes himself to have a reason¹ for believing p; that is, only if there is some proposition or fact that she is prepared to cite in defence of her believing p. Thus, according to Ginsborg, a subject's belief that p can rationalise, or make rationally intelligible, another belief that q only if we can assume that the subject takes p as a reason¹ for believing q.

One point to note here, however, is that a subject's having a reason² for believing p need not entail that a subject thereby has justification for believing p. As I explained earlier, for a subject to be justified in believing p on the basis of his having a reason r (that is, a reason¹) for believing p, it seems essential that the subject has a justified belief that r. It seems possible, however, for a subject to have a belief that p which makes rationally intelligible, in some sense, his having a further belief that q without the latter belief's being justified for him. This kind of situation can obtain, for example, when S's belief that p is itself unjustified for him. On the other hand, we can take it that, when a subject has a justified belief that p which

constitutes, in part, his having a further justified belief that q, the belief that p rationalises his belief that q and is, therefore, a reason² for his having that belief.

Now returning to McDowell's view of perceptual justification and the question of what the subjectivist element in it amounts to, it is clear that McDowell does think that having a reason for believing something consists at least in part of being able to cite a proposition in support of what one believes. Thus quoting Wilfrid Sellars, McDowell endorses the view that:

In characterising an episode or a state as that of knowing, we are not giving an empirical description of that episode or state; we are placing it in the logical space of reasons, of justifying and being able to justify what one says.
(quoted in McDowell 1994, xiv)

As it becomes clear here, to McDowell, a subject's belief that p cannot belong to the space of reasons and, therefore, be properly justified for him unless he is able in an appropriate way to provide a justification or a reason for this belief. In the case of a perceptual belief that p, however, which is based, say, on one's having a visual experience in which it appears to one as if p, the Sellarsian condition, or so McDowell thinks, can be quite straightforwardly satisfied. Thus McDowell claims that "a proper move in the game of giving reasons" can simply be to claim that one sees that p. This, he continues, is "a move that, if one can make it truly, vindicates one's entitlement to a claim with the content of the embedded proposition" (McDowell, *Knowledge and the Internal Revisited*, 2002, p. 100).

It would seem, therefore, that, according to McDowell, when one has perceptual justification for believing p and can, as such, properly justify that belief on the basis

of claiming that one sees that p, one possesses a reason qua proposition (or a reason¹) in the sense outlined earlier. This would mean that one's perceptually justified belief that p involves one's having a justified belief that one sees that p or, at the very least, one's being in a position to have such a justified belief.

Now I think this is true as far as it goes, but we need to be careful here about what exactly, on McDowell's picture, constitutes one's justification for believing p when this involves the ability to appropriately cite the fact that one sees that p as a reason for believing that p. In particular, a correct interpretation of McDowell's views here would need to take account of what he says in the following passage:

Of course one does not inherit entitlement to, for instance, 'There is a candle in front of me' from a commitment – to which one would have to be entitled – to "I see that there's a candle in front of me." One could not be entitled to "I see that there's a candle in front of me" while it was still in suspense whether one was entitled to "There is a candle in front of me" – suspense that one would terminate, on this impossible picture, by inferring "There's a candle in front of me" from "I seen that there's a candle in front of me". But the impossibility of this picture does not disqualify "I see that..." from its status as the form of a proper move in the game of giving reasons, a move that, if one can make it truly, vindicates one's entitlement to a claim with the content of the embedded proposition. The point just brings out the insufficiency of a conception of justification that limits itself to inferential inheritance of entitlement (perhaps with a special story about one's entitlement to the premises of the envisaged inferences). (ibid, p. 100)

Thus, although one's having perceptual (visual) justification for believing *p* involves the capacity to appropriately justify one's belief by citing the fact or proposition that one sees that *p*, one's justification for believing cannot be (solely) constituted by this capacity. The reason for this, so McDowell claims, is that one could not appropriately cite the fact that one sees that *p* in defence of one's believing *p* unless one already had some justification for believing *p*, which does not consist of one's being able to appropriately cite the fact that one sees that *p*.

This essentially brings us back to the passage cited on p.82. There McDowell made clear that when one has justification for believing *p* on the basis of one's having a perceptual experience which amounts to an openness to the fact that *p* (or the fact that makes the proposition *p* true), it is this fact – the fact that *p* – rather than the fact that one sees that *p* that functions as one's reason for believing *p*. In fact, on McDowell's picture, it is having a visual experience of this kind which provides one with justification for believing both the proposition *p* and the proposition that one sees that *p*. Thus he states:

What does entitle one to claim that one is perceiving that things are thus and so, when one is so entitled? The fact that one is perceiving that things are thus and so. That is a kind of fact whose obtaining our self-consciously possessed perceptual capacities enable us to recognize on suitable occasions, just as they enable us to recognize such facts as that there are red cubes in front of us, and all the more complex types of environmental facts that our powers to perceive things put at our disposal. (McDowell, *The Disjunctivist Conception of Experience as Material for a Transcendental Argument*, 2008, p. 387)

In reconstructing McDowell's view, we therefore need to distinguish two senses in which a subject can have a perceptual reason for believing a proposition *p* about his environment. Thus, on the one hand, his having a perceptual (visual) reason for believing *p* consists in his ability to appropriately justify this belief by citing the fact that he sees that *p*. On the other hand, however, his having a perceptual (visual) reason for believing *p* will consist of his having a visual experience which is open to the fact that *p* (or the fact that makes *p* true), a fact which, insofar as it is being experienced by him, can function as his reason for believing *p*.

Now we can assume that the notion of a reason that figures in the first of these two senses of having a reason is that of Ginsborg's notion of a reason¹. The question, however, is how we should understand the notion of a reason that figures in the second of the two senses of having a reason for believing *p*. If the latter notion of a reason is not that of a reason¹, then we cannot take this notion to pick out a proposition or fact that one would cite in an appropriate defence of what one believes. But then what other relevant sense of a reason could be involved here?

Before answering this question, we should note that the issue of how to understand the notion of a reason involved in the second sense in which, on McDowell's view, a subject can have a perceptual reason for believing a proposition *p* about his environment bears on the question of whether or not perceptual experiences can amount to a reason² for believing *p*. Thus recall that a reason², on Ginsborg's understanding of this notion, is a mental state *M* of a subject *S* on basis of which, from a third-person perspective, we can determine whether a given belief *B* of *S* is rationally held by *S*. In particular, if, from a third-person perspective, *M* makes *B* rationally intelligible, then we might take it, as a third-person observer, that

a fact presents itself to S in virtue of his having M which functions as his reason¹ for having B. The question now is whether a fact can present itself to a subject in virtue of having a particular kind of perceptual experience in such way that we can take that experience to constitute a reason² for believing a proposition p about his environment.

The answer to this question may help to settle a dispute between McDowell on the one hand and Donald Davidson on the other. Thus Davidson famously claims “...nothing can count as a reason for holding a belief except another belief.” (Davidson, 2000, p. 416) Ginsborg understands this to amount to the claim that only beliefs can function as reason² for further beliefs; that is, it is only by appeal to a subject’s beliefs that, from a third-person perspective, we can render rationally intelligible a subject’s belief that p. In contrast to Davidson, McDowell does seem to hold that the scope of reasons² should include not just beliefs but also perceptual experiences. This is a claim, however, that Ginsborg seeks to undermine.

If I see that p, but without believing that p, do I nonetheless take p to be a reason¹ for believing that p? It is important to be clear that the question is not just whether I take p to imply q. I take the proposition that the moon is made of cheese to imply the proposition that the moon is edible. But it does not follow that I take moon’s being made of cheese to be a reason¹ for believing that the moon is edible, or in other words that I take myself to have, in the fact that the moon is made of cheese, a reason¹ for believing that it is edible. So for any propositions p and q, there is more involved in taking p as a reason for believing that q than simply taking one to imply the other. Now intuitively it appears that the reason that I do not take the moon’s being made

of cheese as a reason¹ for believing that the moon is edible is that I do not believe that the moon is made of cheese. Taking something as a reason¹ for a belief is a matter of being prepared to cite it in defence of the belief, as a consideration which counts in its favour; and I cannot cite the moon's being made of cheese in defence of anything unless I believe that the moon is made of cheese. But if I am right about the relation between reasons¹ and reasons², this intuition is just what McDowell has to deny in extending the scope of reasons² to include experiences as well as beliefs. He has to be able to say that even if I do not believe that an object presented to me is square, its veridically appearing to me as square can nonetheless make it the case that I regard its squareness as a consideration counting in favour of a belief, in particular the belief that it is square. (Ginsborg, 2006, p. 304)

Here Ginsborg reiterates the claim that if one has a reason¹ *r* for believing a proposition *p* then one has to take *r* as a reason¹ for believing *p*, that is, as a fact or consideration that one would cite in defence of believing *p*. One cannot, however, take a fact or proposition *r* as a reason¹ for believing *p* if one does not believe that *r* obtains. What follows from this, so Ginsborg seems to argue, is, just like Davidson claims, that only a belief can function as a reason² for another belief.

The line of argument that Ginsborg offers in this passage, however, crucially depends on the assumption that a mental state *M* can be a reason² for believing a proposition *p* only if the subject takes himself to have a reason¹ for believing a proposition *p*. The problem with this, however, is that Ginsborg does not consider an alternative conception of a (subjective) reason which would make intelligible the idea that in a perceptual experience one can be presented with a fact which can then

function as one's reason – whatever this sense might be – for believing the proposition that this fact makes true. Thus what McDowell needs in order to resist Ginsborg's objection is some understanding of how a perceptual experience can rationalise a belief – act as a reason² for that belief – by presenting the subject with a fact as a reason for believing a proposition *p* without assuming the subject takes that fact as a consideration that he can cite in defence of believing *p*.

Now I think that an alternative conception of a subjective reason that can present itself to a subject when having the relevant kind of perceptual experience can be developed. However, I would suggest that we need to make some specific assumptions about what this kind of experience would need to involve. Thus, for one, it does not seem possible to explain how a perceptual experience can provide one with a reason *r* for believing a proposition *p* – and how this perceptual experience can be a reason² for believing *p* – if all we assume about this perceptual experience is that it is of a certain phenomenological sort. That is, it is not sufficient to think of a reason-providing perceptual experience as one which is best described from the subjective point of view as one in which an object *x* presents itself to one as having a property *F*. Nor is it enough, or so it seems to me, in explaining how a perceptual experience can be reason-providing, to conceive of it as having a phenomenal character – as the naive realist and disjunctivist might say – that is constituted by a the subject's relation to a mind-independent fact.

However, a further key element in McDowell's picture of perceptual justification, which is alluded to in the quoted passage from *Mind and World*, is that perceptual experiences can provide one with reasons qua facts for believing propositions about one's environment only insofar as these experiences are conceptual. In the first instance, this means that for one to have a genuinely reason-

giving perceptual experience is for one to be in possession of certain concepts. In particular, for one to have an experience which amounts to an openness to the fact that p requires that one possesses the concepts needed in order to think the thought that p. Adopting a term from Alva Noe, we can call this the *dependency thesis*.¹⁶

However, I also take McDowell to endorse a somewhat stronger thesis than the dependency thesis. Thus, it is not only that having a genuine reason-giving perceptual experience requires the possession of certain concepts; it is rather that the experience itself is constituted by the ‘actualisation of conceptual capacities’, namely those very same capacities that are involved in thinking or forming perceptual beliefs and judgements. Call this the *actualisation thesis*.

How exactly is the actualisation thesis to be understood? One way of understanding this thesis is as making the claim that one could not have a perceptual experience of a certain phenomenological sort unless one possesses certain relevant concepts and unless these concepts are ‘actualised’ in one’s having the experience. This, at any rate, seems to be the thesis that is suggested by Noe:

Visual experience is experience of objects and features of the environment. There just is no more basic, more neutral way of describing how things perceptually seem to us than that available to us when we describe our experiences as of the kinds of things and properties and events we take to inhabit the world around us...We pick out or individuate experiences by reference to judgments we would make were we to take the experience at face value. We have visual experiences as of geese flying overhead, deer grazing in the meadow, lions hunting down gazelles, and armadillos crossing

¹⁶ (Noe)

the road. To have visual experiences is not to judge that things are some way or other, but it is to represent things as being some way or other. It is to represent things as being a way we can appreciate them as being. That's why experience requires the mastery of concepts. (Noe, p. 7)

Thus, when one has a perceptual experience, for example, as of geese flying overhead or as of deer grazing in the meadow, one essentially has a perceptual experience of a certain phenomenological kind and one could not have an experience of this kind if one were not capable of judging that geese are flying overhead or that deer are grazing in the meadow. The ability to grasp the judgment or thought that, for example, there are deer grazing in the meadow, is a precondition of one's capacity to have the corresponding perceptual experience.

In response to Noe's contention here, we might find it problematic to think that one's having a perceptual experience of a certain phenomenological sort has to be a matter of one's having, and actualising, certain conceptual capacities. For example, we might hold that children and animals can have perceptual experiences of the same phenomenological sort as conceptually sophisticated mature human beings while lacking the requisite conceptual capacities to make relevant judgments about what they are experiencing.

However, without getting drawn too deeply into the question of whether or not conceptual capacities are required in order for one to enjoy a perceptual experience of a certain phenomenological sort, I want to suggest at this point that the importance of McDowell's suggestion that perceptual experiences involve the actualisation of conceptual capacities in fact lie elsewhere. Thus the issue here is not so much – or not only – that of whether or not one could have a perceptual experience of certain

phenomenological sort only if certain conceptual capacities are actualised in one's having the perceptual experience but rather that of whether a perceptual experience can have certain epistemic properties unless we can assume that one's having such an experience involves one's having certain requisite conceptual capacities.

So what we are seeking is a kind of conceptual-experiential state which makes intelligible how a mind-independent fact that *p* can present itself to one as a reason for believing *p*. To see what this kind of state might amount to, consider by way of analogy the kind of a priori justification one can have for believing conceptual truths such as, for example, that all bachelors are male. One way of accounting for the kind of justification one can have for believing the proposition that all bachelors are male is to say that one's grasp or understanding of this proposition is sufficient for one's having justification for believing, and, in fact, for knowing, that it is true. Note, that one's justification here, or so I would suggest, is not merely objectivist but also subjectivist. That is, it is not just that in grasping the conceptual truth that all bachelors are male, one is objectively connected to the truth of this proposition, it is also that in grasping the proposition, its truth becomes *subjectively transparent to one*. In this sense, the proposition that all bachelors are male can be one's reason for believing this proposition to be true. It becomes one's reason for believing it to be true when it is the content of the state one is in when one grasps that proposition to be true. Crucially, however, it seems that this is not the kind of reason that, when one possesses it, one would cite as a consideration in favour of believing that all bachelors are male. At the very least, it is not the kind of reason for believing a proposition *p* that one possesses when, for example, one would cite the fact that the streets are wet as a reason for believing that it has recently rained. Hence it is not a reason¹ in Ginsborg's sense. Furthermore, the kind of justification one has for

believing that all bachelors are male requires one to be in a state, namely the state of grasping or understanding this proposition, which involves the actualisation of specific conceptual capacities. These, of course, are those conceptual capacities that are required to judge and come to believe that all males are bachelors.

Now, by way of analogy, we might suggest then that there is a perceptual or experiential state, say, the state of seeing that *p*, that one can be in, *in which the obtaining of the fact that p (or of the fact that makes p true) becomes subjectively transparent to one*. Like in the case of one's having a priori justification for believing propositions such that all bachelors are male, the justification that one has for believing an external world proposition *p* when the fact that *p* (or the fact that makes *p* true) becomes transparent to one in seeing that *p*, is not merely objectivist but also subjectivist. When one sees the fact that *p*, that fact is one's reason for believing the proposition that *p*. Once again, the kind of reason one has when one sees that *p* is not the kind of reason, when one possesses it, that one would typically cite as a consideration in favour of believing *p*. Hence it is not a reason¹ in Ginsborg's sense.

But what kind of state exactly is the state of seeing that *p*? It is important to emphasise that this state should not be assimilated to the state of grasping conceptual truths such that all bachelors are male. Nor should we think that the sense of reason involved in the notion that the proposition that all bachelors are male can be one's reason for believing that proposition to be true when one grasps it to be true is the kind of reason involved in the notion that the fact that, say, there is a table in front of one, can be one's reason for believing that fact to obtain when one sees that there is a table in front of one. In the end, I think that the state of seeing that *p* or the state of having a perceptual experience 'open' to the fact that *p*, which figures in

McDowell's view of perceptual justification, is simply a sui generis mental state. In presenting his view, it seems to me, that McDowell relies on the expectation that most of us will find the notion that in seeing that p the obtaining of the fact that p becomes transparent to one as being intuitively compelling. To determine whether there is such a state, the suggestion might be that all that one needs to do is to briefly reflect on one's current situation. In looking around multiple facts just seem to present themselves to one. These facts seem to be made manifest in one's visual experience of them.

At this point, it should become clear that the possibility of being in an experiential state in which the obtaining of a mind-independent fact becomes transparent to one crucially depends on perceptual experiences having the phenomenal character that they do. Thus, it seems clear that we could not conceive of how the obtaining of a mind-independent fact that an x is F can become subjectively transparent to one when one sees that an x is F, unless one is in a visual state that from the subject's own introspective point of view is one that is best described as a conscious awareness of an x being there in front of one and as having the property F. It is, therefore, at least in part, in virtue of a perceptual experience having the phenomenal character that it does that one can be in a state, when having this experience, that provides one with justification for believing propositions about one's environment

It is equally clear that the state that one is in when one sees that p is a conceptual state and involves the actualisation of certain conceptual capacities. I leave open the question of whether the actualisation of conceptual capacities is required for the state of seeing an environmental fact that p to have the phenomenal character that it does. What is clear, however, is that a mental state could not be a state in which the

obtaining of a fact that p becomes subjectively transparent to one unless one has the conceptual capacity to form the judgment that p.

With this, we can sum up McDowell's view of perceptual justification and its application to the Basic Principle as follows. As we have seen, the view contains both an objectivist and a subjectivist element. The objectivist element is that when one has a perceptual experience in which it appears to one as if an x is F and forms the belief that an x is F on the basis of one's having this perceptual experience, one's belief that an x is F is prima facie justified for one to the extent that the perceptual experience constitutes an openness to the fact that an x is F and thereby allows one to form the belief that an x is F on the basis of something that entails (or is necessarily correlated with) the truth of the proposition that an x is F. The subjectivist element, in turn, breaks down into two separate elements. Thus, on the one hand, when one has a perceptual experience in which it appears to one as if an x is F and forms the belief that an x is F on the basis of this perceptual experience, one's belief is justified insofar as one has the capacity to properly justify that belief by claiming that one sees that an x is F. On the other hand, however, the subjectivist element in McDowell's view consists of the notion that in having a perceptual experience in which it appears to one as if an x is F, one can be in a conceptual-experiential state in which the obtaining of the fact that an x is F can become transparent to one.

How then should we evaluate this view? One serious problem, of course, is that this view cannot give us a complete explanation of the Basic Principle because it fails to accommodate in any way the new evil demon intuition. For our evil demon victim can neither satisfy the objectivist condition nor the subjectivist condition in McDowell's view of perceptual justification. He cannot satisfy the objectivist condition because his perceptual experiences do not entail the truth of propositions

he believes on the basis of these experiences, and he cannot satisfy the subjectivist condition because he cannot be in a conceptual-experiential state in which the obtaining of the facts is transparent to him, since this is a factive state. Moreover, McDowell's view not only fails to accommodate the new evil demon intuition, but it also fails to explain how a 'normal' subject, who is not a demon victim, but who inhabits the kind of world that we take ourselves to be in, could have justification for believing a proposition *p* about his environment on the basis of a perceptual experience when, unbeknownst to him, this experience is merely a very deceptive kind of hallucination.

E-Representing the Obtaining of a Mind-Independent Fact

So McDowell's view, as stated, cannot be fully accepted. At this point, he might, of course, offer a two-sense conception of perceptual justification and two-sense explanation of the Basic Principle. Thus, he might say that while in the veridical case one can be justified in believing an external world proposition *p* on the basis of one's being in a state in which the obtaining of the fact that *p* becomes transparent to one, one essentially has a different kind of justification for one's perceptual beliefs when one unwittingly suffers, say, from a highly deceptive hallucination. But this would require some understanding of how one can be justified in one's perceptual belief that *p* when one is not, or is not in a position to be, in a mental state in which the obtaining of the fact that *p* is transparent to one. In the next chapter, I will address the question of whether there is any good reason for adopting a two-sense conception of perceptual justification in the way that a McDowellian might propose. For the

remainder of this chapter, I want to develop a conception of perceptual justification that can explain how one can have a justified perceptual belief independently of whether one is a 'normal' subject or a victim of an evil demon.

As I have suggested earlier, the discussion so far of McDowell's view of perceptual justification can help illuminate an alternative conception of perceptual justification and an alternative explanation of the Basic Principle, both of which can fully accommodate the new evil demon intuition. The problem with McDowell's view, of course, is that, even though it conceives of one's having perceptual justification for believing a proposition *p* about one's environment in terms of one's occupying a certain subjective perspective from which a situation obtains which entails or makes true the proposition *p*, the mental state in virtue of which one occupies this perspective is essentially a factive one. Thus, the experiential-conceptual state that one is in when the fact that *p* becomes perceptually transparent to one is one that entails, or is necessarily correlated with, the fact that *p*. The key to providing an account of perceptual justification and an explanation of the Basic Principle, which is consistent with the new evil demon intuition and with the thought that one can have justification for one's perceptual beliefs even if one suffers from certain kinds of hallucination, is thus to try to conceive of an epistemically significant experiential state, which can constitute a subjective perspective from which a situation obtains that entails or makes true certain external world propositions but which is, crucially, non-factive. One suggestion is that instead of conceiving of perceptual justification as a matter of one's being in a state in which the obtaining of a mind-independent fact becomes perceptually *transparent* to one, we should conceive of perceptual justification as a matter of one's being in a state which perceptually *represents in some way* the obtaining of a mind-independent fact.

What is important here is that the notion of representation in general is not (or not essentially) factive, since one can falsely represent the obtaining of the fact that p.

The question, of course, is how exactly we should conceive of this kind of representational state. As with the state in which the obtaining of the fact that p becomes transparent to one, I suggest that we should think of the relevant kind of perceptual-representational state, if it exists, as in some sense being *sui generis*. What follows is a list of some of the key properties of this kind of state.

First of all, the state is such that, if one is in it, it is liable to fix one's beliefs about the external world. Thus, if one is in a state in which one perceptually represents, in the relevant way, the obtaining of the fact that p, then, in the absence of reasons for believing otherwise, one will typically form the belief that p. However, it is absolutely crucial to emphasize that there is in the subject not merely a disposition to form the belief that p when he perceptually represents, in the relevant way, the obtaining of the fact that p. For what is required here is not simply a conception of a representational state which is liable to cause certain beliefs about the environment but rather a representational state which constitutes an *epistemically significant subjective perspective* from which a situation obtains that entails or makes true the contents of these beliefs. So I suggest that another feature of this kind of perceptual-representational state is that it is not just liable to fix certain beliefs about one's environment but also that it is such that if one were to reflect on the question of why one should think that these beliefs are true, one would be disposed to think – as a result of one's having a perceptual experience of a certain phenomenological sort – that these beliefs are true because one can see their contents as being true or because the facts that make these beliefs true (visibly, say) obtain before one. Thus the state we are concerned with is a perceptual-representational

state of a certain phenomenological sort such that if the subject were to reflect on his current situation when being in this state is he would, in the absence of reason for thinking otherwise, be disposed to think that it is a situation in which certain facts are simply made perceptually manifest to him.¹⁷

It is in these ways that this is a kind of experiential or perceptual state in which one represents the obtaining of the fact that p in an epistemically significant way. I will refer to this kind of state as a *perceptual e-representation* of the mind-independent fact that p. Thus, on the resulting account of perceptual justification, one has justification for believing that p on the basis of one's having a perceptual experience of the relevant phenomenological sort if one perceptually e-represents the obtaining of the fact that p (or the fact that makes true the proposition p).

It should be clear from the forgoing that the phenomenal character of perceptual experience plays an absolutely vital role in the conception of a perceptual e-representation. Thus one could not perceptually e-represent the obtaining of the fact that P unless one enjoys a perceptual experience in which it phenomenally appears to one as if p. We should also emphasize, however, that the state of perceptually e-representing the obtaining of the fact that p is essentially conceptual and as such requires the actualisation of the relevant conceptual capacities. Thus one could not be in this kind of epistemically significant state unless one could form the judgment that proposition p is true.

To further clarify the conception of a perceptual e-representation, it is important to note that the kind of representational state that one is in when one

¹⁷ Compare with this with Martin's notion of the non-neutrality of perceptual experience. See (Martin, *The Transparency of Experience*, 2002, pp. 389-91) However, it is not clear if Martin thinks that perceptual experiences are non-neutral in virtue simply of their having the phenomenal character that they do. That thought, I think, would be mistaken. For it seems to me that perceptual experiences are non-neutral only to the extent that they involve the actualisation of conceptual capacities.

perceptually e-represents the obtaining of a mind-independent fact is not necessarily the kind of representational state that the intentionalist about perceptual experiences claims is the kind of state one is in when one has a perceptual experience of a certain phenomenological sort. Thus, an intentionalist makes a commitment about what kind of state can constitute one's being in a state that has a specific phenomenal character. Thus, he will say, for example, that when one has a perceptual experience in which it appears to one as if an x is F , then one is essentially in a state with a certain representational content which represents an x as being F . However, as I understand it, the conception of a perceptual e-representation of a mind-independent fact that p is entirely metaphysically neutral. Thus, while having a perceptual e-representation requires that the subject is in a perceptual state of a certain phenomenological sort, the conception of a perceptual e-representation is in itself consistent with the various theories about what constitutes the phenomenal character of perceptual experiences. Of course, if intentionalism about perceptual experiences turns out to be the correct account of the phenomenal character of perceptual experiences, we would want to say that if one perceptually e-represents the obtaining of the fact that p that one's being in this state is a matter of, or is constituted by, one's being in a representational state with a certain representational content. Alternatively, however, we might hold on a sense-datum theory that one perceptually e-represents p the obtaining of the fact that p , in part, when having a perceptual experience whose phenomenal character is constituted by one's standing in some relation of awareness to a mind-dependent sense-datum and its properties. Finally, on a disjunctivist conception of perceptual experience, we might hold that one perceptually e-represents the obtaining of the fact that p either when having a perceptual experience whose phenomenal character is constituted by a relation to a mind-independent fact or when having a perceptual

experience whose phenomenal character is a matter of this experience being subjectively indistinguishable from a certain kind of veridical experience.

Finally we should note that contrary to the McDowellian conception of a state in which the obtaining of a fact that *p* is transparent to one, the conception of a perceptual e-representation is not a conception of a state that amounts to one's possession of a reason for believing *p*. Thus there is on the latter conception no actual fact that one perceives that can function as one's subjective reason for believing a proposition about one's environment. I do think, however, that, with some care, we can conceive of a perceptual e-representation as a reason² for a belief. For recall that, according to Ginsborg, a reason² for a belief *b* is a mental state *M* on the basis of which, from a third-person perspective, we can make rationally intelligible a person's having *B*. In particular, Ginsborg's suggestion is that a reason² is a mental state *M* in which things present themselves to a subject in such a way that his having *B* makes rational sense in light of his having *M*. Perceptual e-representations do meet this definition of a reason².¹⁸ At the same time we should note that although we can conceive of a perceptual e-representation as a reason², that the notion of a perceptual e-representation is ultimately a technical notion. Thus, in ordinary circumstances, a third-person observer unfamiliar with the theory of perceptual justification being developed here (that is, nearly everyone in the world) would not explicitly attempt to make rational sense of a person's belief by appeal to that person's having a perceptual e-representation. However, we can think of a subject, in ordinary circumstances, making rational sense of someone else's belief that *p* by assuming correctly, for example, that this person can see that *p* or that this person has a perceptual experience in which it appears to him as if *p* is the case. My

¹⁸ Of course, they fail to be reasons² if we also accept Ginsborg's stipulation that a subject's having a reason² for believing *p* involves his taking it that he has a reason¹ for believing *p*.

suggestion would be that in this case the subject's assumption is correct precisely because the person he is trying to make rational sense of enjoys a state in which he perceptually e-represents the obtaining of the fact that p.

Chapter 3

Introduction

In the last chapter we developed an explanation of the Basic Principle of Perceptual Justification that could fully accommodate the new evil demon intuition and thereby account for the fact that the victims of an evil demon, whose perceptual experiences are generally unreliable but otherwise phenomenologically like our experience (or if one prefers: subjectively indistinguishable from our experiences), could nonetheless be epistemically justified in holding their perceptual beliefs. As we have seen, this explanation of the Basic Principle relies in part on facts about the phenomenal character of perceptual experiences. Thus, if the explanation is correct, then it is in virtue of their phenomenal character that perceptual experiences contribute to the epistemic justification of our perceptual beliefs. Crucially, however, the resulting picture of perceptual justification, while appealing to facts about the phenomenal character of perceptual experience, is consistent with the various theories of the metaphysics of perceptual experience that are on offer.

In this chapter, I want to address the question of whether there is more to say about the relation between the metaphysics and epistemology of perceptual experience. In particular, I will be concerned with whether a disjunctivist conception of perceptual experience can have any further implications on how our perceptual beliefs can be justified for us.

Thus suppose we accept a disjunctivist conception as well as the explanation of the Basic Principle that was developed in the last chapter. According to the latter,

one has justification for believing a proposition *p* about one's environment on the basis of a perceptual experience if, partly in virtue of the given experience and its phenomenal character, one perceptually e-represents the obtaining of the fact that *p* (or the fact that makes *p* true). If this view of perceptual justification is combined with the disjunctivist conception of perceptual experience, we would say that one can perceptually e-represent the obtaining of the fact that *p* either when one has a perceptual experience whose phenomenal character is constituted by one's standing in some relation of awareness to the fact that *p* or when one has a perceptual (hallucinatory) experience whose phenomenal character is simply a matter of this experience's being subjectively indistinguishable from a veridical experience of the fact that *p*. In both the veridical and the hallucinatory case, the kind of justification one has for believing *p*, however, is essentially the same. For in both cases, the justification one has for believing *p* is constituted by one's perceptually e-representing the obtaining of the fact that *p*.

Still, in accepting a disjunctivist conception of perceptual experience, the fact remains that, in the veridical case, one has a perceptual experience whose phenomenal character objectively connects one with the truth of certain propositions about one's environment. So, on this conception, there is potential room to hold that, in the veridical case, one could have a kind of justification for believing a proposition about one's environment that one could not have if one suffered from a hallucination. In particular, we might hold, adopting the McDowellian picture of perceptual justification we considered in the previous chapter, that when one has a veridical experience, given the way the phenomenal character of this experience is constituted, one would not merely perceptually e-represent the obtaining of the fact that *p*, but rather enjoy a state in which the obtaining of the fact that *p* becomes

transparent to one.¹⁹ This in turn could lead us to accept a two-sense conception of perceptual justification (and a two-sense explanation of the Basic Principle) on which one will have different kinds of perceptual justification for believing propositions about one's environment depending on whether or not one has a veridical perceptual experience.

The crucial question, of course, is whether there are any good arguments for insisting that in the veridical case one has a kind of perceptual justification for believing propositions about one's environment which is not available in the hallucinatory case and which is a matter of one's having a perceptual experience that is constituted by one's standing in a relation to a mind-independent fact. To answer this question we need to turn, once again, to some of McDowell's writings. Thus, apart from the claim that perceptual experiences provide one with reasons for believing propositions about one's environment, another key claim in his epistemology is that perceptual experiences can in some cases provide one with a kind of justification which amounts to one's having *conclusive warrants* for external world beliefs. Crucially, McDowell holds that it is only insofar as perceptual experiences can provide one with justification for believing propositions about one's environment which amounts to one's having conclusive warrants that these experiences can provide one with *knowledge* of those propositions. Furthermore, he seems to hold that perceptual experiences can provide one with conclusive warrants for believing external world propositions only if these experiences, in the veridical cases, are constituted by one's standing in some relation to objective facts. This

¹⁹ One would need to explain, of course, how a disjunctivist could secure this McDowellian conception of a reason-giving experience. The question is also whether only disjunctivism could make sense of how such an experience is possible. Finally, we might ask whether disjunctivism could yield a different conception of perceptual justification other than McDowell's conception. I am not aware of any alternative subjectivist conception. Disjunctivism of course could yield a straightforward objectivist conception.

means in turn that in the veridical case, perceptual experiences provide one with a kind of justification that is sufficient, qua justification, for one's coming to know a proposition about one's environment that cannot be had if one were in a hallucinatory case.

Now in evaluating this line of reasoning, we need to have an understanding of the notion of a conclusive warrant and of why having such a warrant is essential to one's having perceptual knowledge of one's environment. We should also ask whether a conclusive warrant cannot be something that one has in both the veridical and the hallucinatory cases and, in particular, whether it is something that one can have even if one's perceptual justification for believing a proposition about one's environment consists of one's perceptually e-representing the obtaining of an environmental fact. But suppose for the moment that the notion of one's perceptually e-representing the obtaining of an environment fact is not in fact a notion of perceptual justification that involves one's having a conclusive warrant for believing an external world proposition, however the notion of a conclusive warrant is to be understood. As such, one's having this kind of perceptual justification, according to McDowell, would essentially fail to support one's having any perceptual knowledge of one's environment. And that, of course, would be an unfortunate consequence if we were to insist that one's having any perceptual justification for believing a proposition *p* about one's environment could consist in nothing else but one's perceptually e-representing the obtaining of the fact that *p*. Thus while we might happily accept that, in the hallucinatory case, one can merely have a kind of perceptual justification for believing *p* which essentially fails to deliver perceptual knowledge of the fact that *p*, it is surely unacceptable, in the veridical case, that one essentially lacks a kind of perceptual justification for believing *p* which, qua

justification, is sufficient for one's having perceptual knowledge that p. Thus, even if, in the hallucinatory case, one may well be justified in having a perceptual belief as a result of one's perceptually e-representing the obtaining of the fact that p, we would be compelled to hold that one can have a different kind of perceptual justification for believing an external world proposition p in the veridical case, which is sufficient, qua justification, to yield knowledge of the fact that p. We would therefore have strong motivation for a two-sense conception of perceptual justification (and a two-sense explanation of the Basic Principle) and, perhaps in turn, some motivation for endorsing a disjunctivist conception of perceptual experience, should this conception turn out to make possible one's having a kind of perceptual justification for believing a proposition p about one's environment that is sufficient, qua justification, to yield perceptual knowledge that p.

Now later in the chapter, I will consider in more detail McDowell's argument for the claim that perceptual knowledge requires the possession on the part of the subject of conclusive warrants for believing propositions about the external world, the extent to which the notion of having conclusive perceptual warrants entails that one cannot have such warrants in both the veridical and hallucinatory cases, and the question of whether the notion of having conclusive perceptual warrants presupposes any particular metaphysical conception of perceptual experiences. Before doing so, however, I first want to turn to a related but different epistemological argument for the disjunctivist conception of experience that McDowell has offered in a recent paper. His claim in this paper is that this conception helps to solve the familiar problem of scepticism about the external world.

McDowell's Response to Scepticism about the External World

It will help to begin with McDowell's own understanding of the disjunctivist conception of perceptual experience. Thus, as he understands it, disjunctivism is a conception of perceptual experience that preserves the idea that in non-deceptive veridical perceptions there can be 'direct perceptual access to objective facts about the environment.' (McDowell, 2008, p. 4) It does so by invalidating the inference from the possibility of certain kinds of deceptive perceptual experience, i.e. certain kinds of hallucination (or illusion), to the conclusion that non-deceptive veridical perceptions, i.e. experiences involving no hallucinatory (or illusory) elements, cannot be constituted by direct perceptual access to objective facts. The relevant deceptive experiences that supposedly support this inference are those that are subjectively indistinguishable from non-deceptive veridical perceptions. The idea is thus that when one considers a non-deceptive veridical perception, say, a visual perception of a red cube, one's 'experience could be just as it is, in all respects, even if there were no red cube in front of [one]' (ibid, p. 378). From this it follows that a non-deceptive visual perception of a red cube cannot be constituted by a mind-independent fact – *that there is a red cube* – being made perceptually accessible to one. Against this, disjunctivism, as McDowell conceives of it, holds that perceptual experiences that are subjectively indistinguishable from non-deceptive veridical perceptions are either, in the non-deceptive veridical cases, perceptual experiences in which an objective fact is perceptually accessible or, in the deceptive cases, perceptual experiences in which it merely appears as if an objective fact is perceptually accessible when actually it is not. Thus the disjunctivist idea that the class of experiences subjectively indistinguishable from non-deceptive veridical

perceptions comprises perceptual experiences that both do and do not involve genuine perceptual relations to objective facts can block the inference from the possibility of deceptive experiences subjectively indistinguishable from non-deceptive veridical perceptions to the conclusion that non-deceptive veridical perceptions cannot involve genuine perceptual relations to objective facts.²⁰

What bearing does the disjunctivist conception of experience have on scepticism about the external world? According to McDowell, it is the following:

... scepticism expresses an inability to make sense of the idea of direct perceptual access to objective facts about the environment. What generates this scepticism is the thought that even in the best possible case, the most that perceptual experience can yield falls short of a subject's having an environmental state of affairs directly available to her. Consider situations in which a subject seems to see that, say, there is a red cube in front of her. The idea is that even if we focus on the best possible case, her experience could be just as it is, in all respects, even if there were no red cube in front of her. This seems to reveal that perceptual experience provides at best inconclusive warrants for claims about the environment. And that seems incompatible with supposing we ever, strictly speaking, know anything about our objective surroundings. The familiar sceptical scenarios — Descartes's demon, the scientist with our brains in his vat, the suggestion that all our apparent experience might be a dream — are only ways to make this supposed predicament vivid. (ibid, pp. 378-8)

²⁰ Note that disjunctivism as formulated in this paragraph makes no explicit commitment about what constitutes the phenomenal character of experiences. However, this fact won't matter for the purposes of this chapter.

Now it is not clear why McDowell claims that scepticism ‘expresses an inability to make sense of the idea of direct perceptual access to objective facts’ for that *idea* seems clear enough. What I will take McDowell to mean here is simply that the sceptic cannot *accept* that perceptual experiences could be constituted by a direct perceptual access to objective facts about the environment. Presumably, this is because of the kind of inference from the possibility of certain kinds of deceptive experience to the nature of perceptual experiences in general that was outlined earlier. So to the sceptic, as McDowell sees it, a non-deceptive veridical experience of an *x* can be just what it is, in all respects *qua* experience, even if there is no *x* in one’s perceptual field. From this the sceptic then apparently concludes that a perceptual experience, even if non-deceptive and veridical, can at best provide one with inconclusive warrant for beliefs or claims concerning one’s environment. And such a warrant is insufficient, *qua* justification, to constitute the possibility of perceptual knowledge.²¹ Hence perceptual knowledge is impossible and scepticism about the external world, at least insofar as one thinks that perceptual experiences are the source of knowledge about the external world, immediately follows. Finally, the familiar Cartesian scenarios, like the evil demon thought experiment, are only supposed to make vivid the idea that perceptual experiences cannot be constituted by a relation to mind-independent facts.

Having outlined the sceptical argument in this way, McDowell offers the following response:

...it constitutes a response [to this kind of scepticism] if we can find a way to insist that we can make sense of the idea of direct perceptual access to

²¹ I will assume here that for one to know that *p* is for one to have a true justified belief that *p* plus whatever conditions are required to rule out one’s being in a Gettier situation.

objective facts about the environment. That contradicts the claim that what perceptual experience yields even in the best possible case must fall short of having an environmental fact directly available to one. And without that thought, this scepticism loses its supposed basis and falls to the ground. (ibid, p. 379)

Thus if, despite the subjective indistinguishability between veridical perceptions and certain kinds of hallucinations, we can insist that perceptual experiences can consist of direct perceptual relations to objective facts about the environment, we no longer have to conceive of the kind of justification that perceptual experiences can provide as insufficient, *qua* justification, for constituting the possibility of perceptual knowledge. Hence if disjunctivism is true, and if one can thereby accommodate the fact that certain deceptive experiences can be subjectively indistinguishable from non-deceptive veridical perceptions without having to deny the possibility that some perceptual experiences can consist of direct perceptual access to objective facts, scepticism ‘falls to the ground’.

Both McDowell’s account of the reasoning leading to scepticism and his response to it raise some pressing questions. First why should we think that perceptual experiences, when they are not conceived of as involving direct perceptual relations to objective facts, at best provide *inconclusive* warrants for beliefs or claims about one’s environment? Second, even if such perceptual experiences at best provide inconclusive warrants for perceptual beliefs, why is having such a warrant insufficient, *qua* having justification, for constituting the possibility of knowledge about one’s environment?²² However, we can, for the

²² Of course, we need to understand what conclusive warrants are. I will turn to this in the next

moment, leave these questions aside. For now, the more pressing issue is whether McDowell is right to contend that the sceptic makes any assumptions at all about the nature of perceptual experience and about whether perceptual experience can ever be constituted by a relation to mind-independent facts. If it can be shown that the sceptic does not make any assumptions at all about the nature of perceptual experience, then McDowell's anti-sceptical response fails and cannot, therefore, be used in defence of the disjunctivist conception of perceptual experience.

Wright's Response to McDowell's Anti-Scepticism

As we have seen, the effectiveness of McDowell's response to scepticism about the external world clearly depends on the claim that the sceptic cannot make sense of how perceptual experiences can provide one with justification for believing propositions about the external world that is sufficient, qua justification, for the possibility of knowledge about one's environment *only* because he cannot accept that perceptual experience can consist of a perceptual relation to an objective fact about one's environment. According to McDowell, the phenomenological fact concerning the subjective indistinguishability between non-deceptive veridical perceptions and certain kinds of hallucination and the sceptical Cartesian scenarios that make this fact so vivid thus blind the sceptic from a conception of perceptual justification that, if one were not misled by the phenomenological fact, would seem to be sheer common sense. In response to McDowell, Crispin Wright has recently argued that

section. For the moment, I will understand the notion of one's having conclusive warrant as a particular form of one's having epistemic justification.

this reconstruction of the line of reasoning leading to Cartesian scepticism about the external world misconstrues the true import of the Cartesian sceptical scenarios and the subjective indistinguishability between veridical perceptions and certain kinds of hallucinations. For properly understood, the sceptical line of reasoning does not depend at all on any particular conception of what constitutes the experiential nature of perceptual experiences. In illustrating this point, Wright appeals to the following case:

Suppose I knowingly participate in double-blind trials of a new hallucinogenic drug. Half the participants receive the drug, the other half an identical looking and tasting vitamin pill. The consciousness of those who ingest the drug will move — so it is predicted — seamlessly into a completely plausible, sustained, multi-sense hallucination, in which all the experiences “make sense” in the context of the subjects’ recollection of their previous waking experiences of the day prior to ingestion of the pill, and indeed of the general fabric of their lives. This will happen — if the drug performs as expected — within a few minutes of ingestion, and the hallucination will then be sustained for several hours. (Wright, Comments on John McDowell's 'The Disjunctive Conception of Experience as Material for a Transcendental Argument', 2008, pp. 398-9)

This case unfolds as follows. Not knowing which of the two pills it is, I ingest the vitamin pill and continue to have non-deceptive, fully veridical perceptual experiences. Thus, according to the disjunctivist conception, I continue to have perceptual experiences which consist of direct perceptual relations to objective facts.

But that clearly does not entail that I continue to have perceptual justification for any beliefs about those facts. As Wright argues:

...it would, in the circumstances, be unwarranted for me to believe any of the mundane propositions of whose truth-conditions I am perceiving the worldly satisfaction. I do not know that these propositions are true. I am not, in the circumstances, even justified in claiming that they are true (the trialists, remember, are knowingly divided 50-50 but blind between those who take the drug and those who take the vitamins.)...the perceptual apprehension of a state of affairs that makes it the case that P is one thing, and possession of warrant for taking it to be the case that P is another. The former, even when a subject's belief that P is based on that very apprehension, is insufficient for warrant if aspects of the subject's collateral information conspire to make the belief that P somehow irrational or irresponsible. (ibid, p. 399)

The crucial point this case raises is thus that my knowing participation in the trial puts me in an informational state such that I can no longer rationally avail myself of any justification my perceptual experiences supposedly provide. And this is so however the nature of a perceptual experience is constituted

I take it that there is no question that whatever form of justification for empirical beliefs perceptual experiences can provide, this justification can be undermined in this way. Thus Wright raises a straightforward point about the potential defeasibility of perceptual justification. The question, however, is what bearing this has on the sceptical argument. Here Wright seems to think that while the collateral information that one obtains when participating in the above trial undermines whatever warrant

one's perceptual experiences yield at the time around the trial, the raising of the metaphysical possibility of Cartesian sceptical scenarios represents collateral information that undermines the epistemic credentials of perceptual experiences *as such*. Again, quoting from Wright:

...it is undeniable that if the alternatives are to suppose that my current experiences are elements of veridical perceptual activity and to suppose that they are a marvellously convincing counterfeit, then the subjective quality of the experiences itself — what it is like to undergo them — can indeed provide no rational motive for either view. (ibid, p. 400)

It becomes clear therefore how, according to Wright, McDowell misconstrues the sceptical line of reasoning. McDowell is right in claiming that the sceptic cannot make sense of how perceptual experiences can provide warrants that are sufficient, *qua* justification, for the possibility of knowledge about one's environment. But this is not because the sceptic cannot accept that perceptual experience can consist of a perceptual relation to an objective fact about one's environment, as McDowell would have it. Rather it is because the sceptic cannot accept that one can rationally avail oneself of whatever form of justification perceptual experiences might provide in light of the metaphysical possibility of sceptical scenarios – for example, that one is brain-in-a-vat – and the phenomenological fact that if one were a brain-in-a-vat, one's experiences would be subjectively indistinguishable from non-deceptive, fully veridical perceptions. So Wright claims that McDowell's response to Cartesian scepticism about the external world is 'dialectically quite ineffectual' (Wright, (Anti-) Sceptics Simple and Subtle: G.E. Moore and John McDowell, 2002, p. 331). It will

not do to simply insist on the disjunctivist conception of perceptual experience and the conception of perceptual justification that it supposedly makes available.

Now to see if Wright is correct in claiming that McDowell's response to scepticism is ineffectual, we need to take a closer look at Wright's alternative reconstruction of the sceptical reasoning. As became clear, the epistemic predicament the sceptic describes is the result of the possession of collateral information concerning the metaphysical possibility of being in sceptical scenarios that threatens to undermine the epistemic credentials of one's perceptual experiences. In driving home this point, Wright relies on the analogy between raising the possibility of one's entrapment in a sceptical scenario and knowingly taking part in the trial of the hallucinogenic drug. But how well does the analogy hold?

As the drug trial case is described, what it represents is a situation in which one is rationally compelled to consider the question of whether one's perceptual faculties are properly functioning and to regard this question as entirely open.²³ The reason for this is that one has evenly balanced evidence for and against the claim that one's perceptual faculties are properly functioning. It is clear, however, that the raising of the metaphysical possibility of being in a sceptical scenario does not create the exact same kind of evidential situation. That is, it is not the case, in the context in which, for example, the brain-in-the-vat hypothesis is raised, that one has evenly balanced evidence for and against the claim that one is not a brain-in-a-vat as a result of having evidence for the claim that there is an exactly even chance that one is a brain-in-a-vat and that one is not. Nevertheless, Wright believes that the effect of raising the brain-in-the-vat hypothesis is to create an epistemic situation that is no better than – or analogous to – one's epistemic situation in the drug trial. Wright claims:

²³ See Wright (2008), p. 400.

‘...a special case of balanced evidence is the situation of *no evidence* either way.’ (Wright, 2008, p. 400) And, quoting again a statement mentioned earlier: ‘...it is undeniable that if the alternatives are to suppose that my current experiences are elements of veridical perceptual activity and to suppose that they are a marvellously convincing counterfeit, then the subjective quality of the experiences itself — what it is like to undergo them — can indeed provide no rational motive for either view.’ (ibid, p. 400) But even if one agrees with Wright that a situation of no evidence either way is just a special case of having balanced evidence, it remains questionable whether there are sufficient similarities between the epistemic situation one is in the drug trial case and the epistemic situation one is in when a sceptical hypothesis is being raised. For there is a crucial difference between the two cases in what supports the assumption that one has balanced evidence for and against the claim that one’s perceptual faculties are properly functioning and, therefore, in what generates the (purported) rational compulsion (a) to consider the question of whether one’s perceptual faculties are functioning properly and (b) to regard this question as entirely open. In the drug trial case, what supports this assumption is one’s knowledge of how the trial is set up; namely that half the participants are given the hallucinogenic drug, while the rest are given the vitamin pill. This knowledge is set against whatever form of justification one’s perceptual experiences might otherwise provide, and it is clear that it is sufficient to undermine the epistemic credentials of one’s experiences. And, the point bears repeating: one’s knowledge about how the trial is set up genuinely undermines the epistemic credentials of one’s perceptual experiences precisely because it legitimately throws up the question of whether one’s perceptual faculties are functioning properly and rationally compels one to regard that question as entirely open. There is, however, no analogous support for the

assumption that one has balanced evidence for and against the claim that one's perceptual faculties are functioning properly in the context in which a sceptical scenario is raised. As Wright's reasoning in the quoted passages makes clear, what supports the assumption that one has balanced evidence either way in this context is what he calls 'the subjective quality' of one's experiences, which is the fact that our perceptual experiences are subjectively indistinguishable from certain kinds of hallucination.

But why should this phenomenological fact about our perceptual experiences entail that we can have no (perceptual) evidence either way for believing that our perceptual experiences are genuinely veridical or certain kinds of hallucination? Now to say that one's perceptual experience is subjectively indistinguishable from a certain kind of hallucination is to say that one cannot know on the basis of introspecting one's experience alone that the experience is indeed a veridical perception and not a very deceptive kind of hallucination. But this fact entails that we can have no evidence either way for believing that our perceptual experiences are genuinely veridical or certain kinds of hallucination only if we make the further assumption that evidence for believing that our perceptual experiences are veridical can only be acquired on the basis of introspection of our perceptual experiences. Thus we can summarise the sceptic's argument, as it is reconstructed by Wright, as follows:

- (1) If I have some justification/evidence/rational motive for believing that I am not a brain-in-vat when having a perceptual experience *e*, I need to be able to tell on the basis of introspecting *e* alone that I am not a brain-in-vat suffering from systematic hallucinations.

(2) I am not able to tell on the basis of introspecting e alone that I am having a veridical perception and not a very deceptive kind of hallucination.

Therefore,

(3) I have no justification/evidence/rational motive for believing that I am not a brain-in-a-vat when having a perceptual experience e .

Furthermore,

(4) If I have no warrant/evidence/rational motive for believing that I am not a brain-in-a-vat when having a perceptual experience e , then I have no justification/evidence/rational motive for believing that ordinary propositions such that I have two hands.

Therefore,

(5) I have no justification/rational motive for believing that I have two hands when having a perceptual experience e .

Premise (2) simply expresses the fact, which we presumably we accept, that our perceptual experiences are subjectively indistinguishable from certain kinds of hallucination. Proposition (3) in this argument unproblematically follows from premises (1) and (2). Premise (4) in turn is derived from the so-called closure principle which we can also take to be unproblematic and which states that:

If S knows that p , and knows that p entails q , then S knows q .

Finally proposition (5) in this argument unproblematically follows from (3) and (4).

At issue, therefore, is the status of premise (1) and the question is why we should follow the sceptic in accepting it. For the obvious alternative is suppose that

one can have some justification for believing that one is not a brain-in-vat and that this justification can be acquired on the basis of sources other than what one can come to know on the basis of introspecting one's perceptual experiences. In this case, the fact that veridical perceptual experiences can be subjectively indistinguishable from very deceptive forms of hallucinations would have no bearing on the question of whether or not one can come to know that one is not in a sceptical scenario. If one accepts that the sceptical reasoning can be undermined in this way, then there seem to be two options available to the anti-sceptic. Thus, on the one hand, we can attempt to show that one can have some form of a priori justification for believing that one is not a brain-in-vat, where this is a kind of justification that derives neither from one's having specific kinds of perceptual experience nor from one's introspection of those experiences. On the other hand, we can attempt to show that one can have justification for believing that one is not a brain-in-vat on the basis of one's having some perceptual justification for believing ordinary propositions about one's environment such that one has two hands – propositions that entail that one is not a brain-in-vat. In this case, one's justification for believing that one is not a brain-in-vat thus derives from one's having a perceptual experience which provides one with justification for believing that one has two hands and from one's grasping that the proposition that one has two hands entails the proposition that one is not a brain-in-vat.

In his own approach to answering the sceptic, Wright himself explores the option of whether we can indeed have some relevant form of a priori justification for accepting that one is not a brain-in-vat.²⁴ But suppose that this strategy essentially fails. What then rules out the claim that one can have justification for believing that

²⁴ (Wright, Warrant for Nothing (and Foundations for Free)?, 2004)

one is not a brain-in-vat on the basis of one's having perceptual justification for believing ordinary propositions about one's environment? This question, of course, is crucial, for suppose McDowell were to insist at this point that it is only because the sceptic is simply not open to the idea that if one has a perceptual experience which is constituted by a relation to mind-independent facts one can have conclusive warrant for believing that one has two hands, that the sceptic fails to understand that one can have justification for believing that one is not a brain-in-vat. Thus, according to this rejoinder on the part of McDowell, the sceptic, contra Wright's suggestion, would in fact have to implicitly assume the impossibility of perceptual experiences that are constituted by a relation to mind-independent facts.

However, even though I think that this is how McDowell should respond to Wright's objections, problems remain with this rejoinder. One of these is that McDowell's anti-sceptical rejoinder, even if ultimately correct, may still leave us unsatisfied. For ultimately, what we may have here is simply a combination two conflicting intuitions. Thus the sceptic – and those sympathetic to his reasoning – may simply insist that given the subjective indistinguishability between the kind of veridical perceptual experiences that we take ourselves to enjoy and the kinds of hallucinatory experiences that the brain-in-vat has to suffer, we cannot in fact have justification for believing that we are not brains-in-vats on the basis of our having perceptual justification for believing ordinary propositions about our environment. The McDowellian in turn might simply insist on the intuition that when we have a perceptual experience which is constituted by a relation to a mind-independent fact we can have a kind of perceptual justification that puts in a position to know ordinary propositions about our environment and therefore puts us in a position to know that we are not brains-in-vats. Of course, insofar as we assume that one cannot

have a priori justification for believing that one is not a brain-in-vat and insofar as we assume that we do have justified perceptual beliefs, then we might left in a situation in which we simply have to resist the sceptic's intuition and endorse the McDowellian intuition instead.

But this leads us to a second problem with McDowell's potential rejoinder to Wright's objections. For it is far from clear that McDowell's conception of how perceptual experiences can provide us with justification for believing proposition about the external world, which relies on the assumption that some perceptual experiences can consist of a relation to mind-independent facts, is the only way to resist the sceptic's intuition. Thus, there may well be other ways of providing an intuitively plausible account of how perceptual experiences can contribute to the justification of beliefs about the external world that do not depend on the disjunctivist assumption that in some cases perceptual experiences can consist of a relation to mind-independent facts. Ultimately, therefore, whether or not we should accept a disjunctivist conception of perceptual experience will depend – from an epistemological point of view at least – on whether or not we should accept McDowell's claim that only one's having conclusive perceptual warrants can ensure that perceptual experiences can provide one with knowledge of the external world. McDowell, it seems to me, needs to defend this claim independently of whether or not it can constitute a response to scepticism about the external world. In fact he does provide independent arguments in defence of this claim and it is to these that I will turn in the next section.

Perceptual Knowledge and Conclusive Warrants

In reconstructing McDowell's response to scepticism, I earlier quoted him as saying that

[what] generates...scepticism is the thought that even in the best possible case, the most that perceptual experience can yield falls short of a subject's having an environmental state of affairs directly available to her...This seems to reveal that perceptual experience provides at best inconclusive warrants for claims about the environment. And that seems incompatible with supposing that we ever, strictly speaking, know anything about our objective surroundings.

Continuing, McDowell claims:

...it constitutes a response [to this kind of scepticism] if we can find a way to insist that we can make sense of the idea of direct perceptual access to objective facts about the environment. That contradicts the claim that what perceptual experience yields even in the best possible case must fall short of having an environmental fact directly available to one.

Now whether or not these remarks constitute a reasonable response to scepticism about the external world, they can be taken as making clear McDowell's commitment to the claim that perceptual experiences can provide one with knowledge about the external world only if they can provide one with justification

for believing propositions about the external world that amounts to one's having conclusive warrants for believing those propositions and for the claim that only if perceptual experiences are constituted by a relation to a mind-independent fact can they provide one with these kinds of conclusive warrant. In this section, we need to determine first whether perceptual knowledge really does require one's possessing conclusive perceptual warrants and, second, if knowledge really does require one's possessing conclusive perceptual warrants, whether one's having such warrants requires one's having perceptual experiences which are constituted by a relation to mind-independent facts. Of course, if perceptual knowledge of an external world proposition *p* is possible even on the assumption that one does not have justification for believing *p* that amounts to one's having conclusive warrants for believing *p*, or if one's having perceptual conclusive warrants for believing *p* does not require one's having a perceptual experience which is constituted by a relation to a mind-independent fact, then we may lose any motivation for adopting the disjunctivist conception of experience and for thinking that there is a kind of perceptual justification that one can have only if one has a veridical perceptual experience.

How then is the claim that perceptual knowledge requires the possession of conclusive warrants to be understood? As we have seen in the previous chapter, to McDowell, knowledge is essentially a standing in the space of reasons, which means that one can know a proposition *p* only if one has some reason for believing *p*. We have also seen that McDowell's epistemological outlook consists of two ways in which we can understand the notion of one's having a perceptual reason for believing *p*. Thus, on the first understanding of this notion, one has a perceptual reason for believing *p* if one can appropriately cite the proposition, say, that one sees that *p* as a reason for believing *p*. Here a perceptual reason consists of a proposition

that one can cite in defence of one's belief. On the second understanding of this notion, however, one has a perceptual reason for believing p if, seeing the fact that p , the obtaining of the fact that p becomes visually transparent to one. On this second way of understanding the notion of having a perceptual reason, one's reason for believing p is the mind-independent fact that makes this proposition true. In what follows, we can take it that both these kinds of perceptual reason can amount to the kinds of conclusive warrant that McDowell alludes to in the passages quoted. Having a perceptual reason for believing p , in either of the two senses, amounts to having a conclusive warrant in the sense that one has a reason for believing p which entails, or necessarily makes true, this proposition. We therefore need to examine first the claim that perceptual knowledge of the proposition p requires one's being able to justify one's believing p by citing a conclusive perceptual reason and address the question of whether one's ability to justify one's believing that p requires that one be in a perceptual state which is constituted by a relation to a mind-independent fact. Second, we need to examine the claim that perceptual knowledge of the external world proposition p requires one's being in a state in which the obtaining of the fact that becomes perceptually transparent to one.

Before we begin our examination of these two claims, however, I want to briefly introduce a term – borrowed from McDowell – for the kind of perceptual experience that we would have on a non-disjunctivist conception of experience. The assumption of an alternative conception of perceptual experience would have to be that a non-deceptive, veridical experience of an x could be just as it is, in all respects *qua* experience, even if there is no x in one's environment. Borrowing a term McDowell uses, I will call experiences conceived of in this way *highest common factor experiences* or *HCF-experiences*. McDowell's claim, therefore, is that HCF-

experiences could at best provide inconclusive warrants for beliefs about one's environment.

Consider then the claim that perceptual knowledge that *p* requires one's ability to appropriately justify one's believing that *p* by citing, as a conclusive reason, the proposition that one sees that *p*. This claim entails that one could not have perceptual knowledge that *p* if one is merely able to justify one's believing that *p* by citing a proposition that does not entail the proposition that *p*. For example, one could not have perceptual knowledge that *p* if all one could say, or if all one were justified in saying, in defence of one's believing that *p* that it perceptually appears to one as if *p* is the case. This kind of proposition about the experiential state that one is in is entirely consistent with the falsity of *p* and, in that sense, an inconclusive warrant for believing *p*.

But why should we think that one's merely being able to cite a nonconclusive reason or warrant for believing *p* is incompatible with knowing that *p*? Adapting a line of thought that McDowell employs in a different context, I will assume here that McDowell would reject the idea that one can know that *p* if one is merely able to cite a non-conclusive reason for believing that *p* because this idea would entail that knowing that *p* "could be constituted by being in a position, which for all one knows, [is compatible with the falsity of *p*]. And that seems straightforwardly incoherent" (McDowell, *Criteria, Defeasibility, and Knowledge*, 1982, p. 457). Thus, if one's reason for believing that *p* is, for all one knows, incapable of ruling out the possibility that *p* is false, then it would be incoherent to ascribe to one knowledge that *p*.

Now, from an initial, pre-theoretical viewpoint, this does seem to be straightforwardly true. Thus, if someone were to provide as a reason for thinking *p*,

say, *that it (perceptually) seems to him as if p*, one might rightly hesitate in attributing to him knowledge that p. But, of course, in the normal run of things when one has no doubts about the accuracy of one's visual experiences, one would not hold back from simply claiming that one sees that p and offer instead the comparatively weaker claim that it visually seems to one as if p. McDowell's point, however, seems to be this: that an HCF-experience could never provide one with justification for believing p by making it appropriate for one to claim in defence of one's believing p that one sees that p.

This latter contention can be questioned, and I will question it below. But suppose one accepts it. In doing so, one need not hold that the weaker kind of perceptual claim would actually be made in the ordinary run of things and thereby, by itself, constitute a proper reason for an empirical belief; rather one would take such a perceptual claim as a starting point in a theoretical account of how making such a perceptual claim could figure in constituting one's justification for beliefs about one's environment. It is precisely theoretical accounts of this kind that are the target of McDowell's arguments:

So we are to try to reconstruct the epistemic satisfactoriness implicit in the idea of seeing that things are thus and so, using the following materials; first, the fact that it looks to a subject as if things are that way; second, whatever further circumstances are relevant (this depends on the third item); third, the fact that the policy or habit of accepting appearances in such circumstances is endorsed by reason, in its critical function, as reliable. (McDowell, *Knowledge and the Internal*, 1995, p. 880)

The specific theoretical account that McDowell addresses here says that one's environmental knowledge that *p* can be constituted by the fact (or rather by one's justifiably believing) that it perceptually seems to one as if *p* and the fact (or one's justifiably believing) that such seemings are a reliable indicator for the truth of *p*. Concerning perceptual justification, this account would say that one's justification for believing *p* would be constituted by one's having justification for believing (a) that it perceptually seems to one as if *p* and (b) that such seemings are a reliable indicator for the truth of *p*. A crucial question with this is, of course, whether one can really have an adequate justification for thinking (b). But as far as I can see, McDowell does not dispute the possibility of having such justification. To him, the fault with this view, as he makes clear in the following passage, lies elsewhere:

...the trouble is this: unless reason can come up with policies or habits that will never lead us astray, there is not enough here to add up to what we were trying to reconstruct. Seeing that things are thus and so is a position one cannot be in if things are not thus and so. Given that one is in that position, it follows that things are thus and so. And if reason cannot find policies or habits that are utterly risk-free, the reconstructing materials cannot duplicate that. However careful one is in basing belief on appearances, if one's method falls short of total freedom from risk of error, the appearance plus the appropriate circumstances for achieving the method cannot ensure that things are as one takes them to be. (ibid, p. 880)

Thus the apparent problem with the view under consideration is that it cannot duplicate the epistemic situation that one would be in if one could legitimately claim

to perceive that things are thus and so. Perceptual claims of this kind are factive and therefore offer conclusive reasons for the propositions they embed. In contrast, what one has on the alternative conception of perceptual justification are reasons – (a) and (b) – that, although perhaps one has justification for believing them, do not strictly speaking entail that *p*. But why does this amount to a problem? Recall that McDowell's contention was not the trivial claim that inconclusive warrants are weaker than conclusive warrants but rather that the inconclusive warrants HCF-experiences can yield are not sufficient, *qua* having justification, for constituting the possibility of knowledge about one's environment. But what reasons has he offered so far for thinking that this should be so? If one could really have justification for taking it that (a) and (b) are true, why would having these reasons, *qua* having justification, not be sufficient for constituting one's knowing *p*?

In the above passage, McDowell alludes to the requirement, when forming a belief about one's environment on the basis of inconclusive reasons such as (a) and (b), for a total freedom from risk of error. Again, the suggestion seems to be that this requirement is needed in order to ensure that when one forms an empirical belief on the basis of (a) and (b), one can be in an epistemic situation that is like the epistemic situation one is in when one has factive reasons. However, this is just a restatement of the claim that knowledge requires one's possessing conclusive reasons. What we need it is some independent argument in support of this claim.

So far, this is thus utterly unconvincing. But McDowell raises a separate, perhaps more serious problem for the view he is considering:

I want to urge another problem about the hybrid conception of knowledge. In the hybrid conception, a satisfactory standing in the space of reasons is only

part of what knowledge is; truth is an extra requirement. So two subjects can be alike in respect of the satisfactoriness of their standing in the space of reasons, although only one of them is a knower, because only in her case is what she takes to be so actually so. But if its being so is external to her operations in the space of reasons, how can it not be outside the reach of her rational powers? And if it is outside the reach of her rational powers, how can its being so be the crucial element in an intelligible conception of her knowing that it is so – what makes the relevant difference between her and the other subject? Its being so is conceived as external to the only thing that is supposed to be epistemologically significant about the knower itself, her satisfactory standing in the space of reasons. That standing is not itself a cognitive purchase on its being so...But then how can the unconnected obtaining of the fact have any intelligible bearing on an epistemic position that the person's standing in the space of reasons is supposed to help constitute? How can it coalesce with that standing to yield a composite story that somehow adds up to the person's being a knower? (ibid, p. 883)

The concern here is still with the view that one's perceptual justification for a belief about one's environment consists in one's having justification for believing propositions (a) and (b), which in turn can be one's reasons for believing an environmental proposition p. McDowell calls such a view a hybrid view because it conceives of knowledge as the combination of one's standing in the space of reasons and the fulfilment of the truth requirement. That is, on this account, one counts as knowing p if one's believing p is based on one's believing with justification propositions (a) and (b) and if p is indeed true. That one's believing that p is a matter

of standing in the space of reasons derives from the fact that it is a rational cognitive achievement; a cognitive state one is in because one is led by reasons. That *p* is true, however, is a mind-independent circumstance, independent of any of one's cognitive achievements and therefore outside the space of reasons. McDowell's preferred view of perceptual knowledge is, presumably, very different. On this view, one counts as knowing *p* if one is able to justify one's believing *p* by appeal to a perceptual claim such that one sees that *p* and such a perceptual claim is justified for one in virtue of the direct visual accessibility to one of the objective fact that *p*.²⁵ There is, on this view, no appeal to a mind-independent circumstance that, by definition, would be outside one's standing in the space of reasons.²⁶

According to McDowell, the hybrid view, if this view includes the claim that knowledge can be a standing in the space of reasons, is incoherent. Thus, his point seems to be that on the hybrid view, one could never achieve a standing in the space of reason in believing *p*. This conclusion is reached by considering two subjects. Both subjects are alike in their (purported) standing in the space of reasons in thinking *p*. That is, both their believing *p* is (purportedly) justified in virtue of their having the same justification for taking it that (a) and (b). *p* is true in one case but not the other. Thus, as per the hybrid view, only one subject knows while the other does not. But, so McDowell's reasoning goes, since both subjects are alike in their (purported) standing of the space of reasons in thinking *p*, and yet only one subject

²⁵ Or as I have suggested in the last chapter in virtue of one's being in a conceptual-experiential in which the obtaining of the fact that *p* is transparent to one

²⁶ Although conditions would have to be built into the notion of visual accessibility to the fact that *p* other than that one has a visual experience that is constituted by a relation to the fact that *p*. For as thought experiments like Goldman's fake-barn scenario show, forming a belief that *p* on the basis of a veridical visual experience, even if this experience is constituted by relation to the fact that *p*, need not be enough to know that *p*. Thus in this case one's having a visual experience which is constituted by a relation to the fact that *p* does not amount to one's having visual access to the fact that *p*. One may wonder, contrary to McDowell, if these further conditions – whatever they are – have to be built into the notion of such visual accessibility to the fact that *p* whether having a belief based on such accessibility is really purely a standing in the space of reasons.

knows in virtue of *p* being true in her case but not in her counterpart's, *p* has to be outside the rational powers of both subjects. This means, I take it, that neither subject can in fact have a reason for thinking *p* is the case. Since *p* is outside the rational powers of the knowing subject, having a reason for believing *p* cannot be part of her standing in the space of reasons. It cannot therefore be part of what constitutes her knowing *p*. So her knowing *p* cannot be a matter of standing in the space of reasons.

However, the line of reasoning here seems to me to be invalid. Thus, McDowell is certainly right in claiming that the hybrid view entails that knowledge cannot be merely a matter of one's standing in the space of reason; that is, it cannot, like on McDowell's alternative view, be constituted simply by one's believing that *p* on the basis of one's having a reason for believing that *p*. But he seems to me to be wrong – or unjustified – in saying that the hybrid view entails that knowledge that *p* cannot be a matter at all of one's standing in the space of reason, that is, a matter of one's believing that *p* on the basis of one's possession of reasons for doing so. Thus, take the following simple example. Suppose, at three o'clock in the morning, I see that the lights in my neighbour's house are on. My neighbour is a freelance journalist and I know that he usually goes to bed by eleven o'clock at night and that if the lights are on at his hour, it is most likely because he's working towards meeting a deadline for submitting a piece. I form the belief that he's working on a piece. Clearly, the evidential situation I am in is such that I can be in the exact same kind of evidential situation and yet not know that neighbour is working on a piece because he is, uncharacteristically, watching the whole final season of *The Sopranos*. But that does not entail that my rational powers cannot reach the fact that my neighbour is working on a piece; that is, it does not entail that I cannot have a reason for thinking that he is. Thus, it seems clear that I can, in this case, have a belief that my neighbour is

working on a piece that is based on my having a reason for believing this proposition to be true. Thus, if as the hybrid view entails I can know that my neighbour is working on a piece as a matter of believing that he is and believing this on the basis of my knowing that the lights in his house are on and that the lights are on usually only when he works towards a deadline, then the hybrid view entails that knowledge can be a matter of a standing in the space of reasons.

McDowell cannot accept the idea that two subjects can be alike in their standing in the space of reasons in thinking *p* and yet only one subject knows, or has the opportunity to know, *p*. It seems to me, however, that we cannot reasonably accept that idea only if we already accept the thesis that only having conclusive reasons amounts to having any reasons at all. But there is, as far as I can see, nothing in McDowell's arguments that amounts to an independent defence of that thesis.

Of course the hybrid conception is deficient for other, more straightforward reasons. Thus, it depends on the assumption that one can have some justification for believing an external world proposition *p* on the basis of one's having some prior justification for believing that it perceptually seems to one as if *p* and for believing that such seemings are a reliable indication that *p* is the case. It is doubtful, to say the least, that one can really have any prior justification for the second of these beliefs.

At this point, however, we might ask why we should accept the assumption that an HCF-conception of perceptual experience cannot support the idea that one's justification for believing an external world proposition *p* can be matter of one's being able to legitimately justify this belief on the basis of citing a conclusive reason of the form that one sees that *p*. Thus, it seems that even if one is suffering from a hallucinatory experience as of *p*, one can appropriately justify one's believing *p* on the basis of claiming that one sees that *p*. We can then say that this kind of claim

would be legitimate under the circumstances because one has a perceptual experience of a certain phenomenological sort and because one possesses the right sort of conceptual capacities. In this way, one's justifying one's belief that p by saying that one sees that p is legitimate because one is in a perceptual state which constitutes a subjective perspective in which fact that p appears to be visibly manifest to one. It is far from clear, therefore, why the insistence that that perceptual knowledge requires the possession on the part of the subject of conclusive reasons or warrants, in the sense we have looked at, requires the disjunctivist conception of experience.

With this, we can turn to the second sense in which we can understand the notion that perceptual knowledge requires conclusive warrants or reasons. Here the claim is that perceptual knowledge that p requires conclusive warrants in the sense that it requires one's being in a conceptual-experiential in which the obtaining of the fact that p (or the fact that makes p true) is perceptually transparent to one. In this case, one's reason for believing p is the fact that p which entails or makes true the proposition p and is therefore conclusive with respect to that proposition. The first question we might ask then is whether one's having this sort of reason for believing an external world proposition p requires that one is in a perceptual state which is constituted by a relation to the mind-independent fact that p. Suppose, for the moment, that this requirement indeed holds. We can then say that one's being in the state in which the obtaining of the fact that p is perceptually transparent to one is a matter of one's being in a perceptual state which is constituted by one's standing in a relation to a mind-independent fact and by one's actualising certain conceptual capacities. However, it is not clear why we could not hold instead, on an alternative conception, that one's being in a state in which the obtaining of the fact that p is

perceptually transparent to one is constituted by one's having a relevant HCF-experience which caused in the right sort of way by the fact that p and by one's actualising certain conceptual capacities. If the latter is indeed coherent, then even if we correctly assume that perceptual knowledge that p requires one's being in a state in which the obtaining of the fact that p is perceptually transparent to one, we no longer have a clear epistemic motivation for endorsing the disjunctivist conception of perceptual experience.

Suppose, however, that we do accept that an HCF-conception of perceptual experience cannot make intelligible how one could be in a conceptual-experiential state in which the obtaining of the fact that p is perceptually transparent to one. In this case, an HCF-conception could not make sense of how one can have the relevant kind of conclusive warrant for believing an environmental proposition p. But should we in this case also accept that perceptual knowledge that p can only be secured if one's justification consists of one's having a fact made manifest to one in this way?

Now McDowell does not offer any further arguments for the claim that one's having conclusive perceptual warrants for believing p is a necessary condition for one's having perceptual knowledge that p. So we have to return to the arguments that we have already considered in connection to the first sense of one's having conclusive warrants. The question is thus whether these arguments fare any better in connection with the second sense of one's having conclusive warrants.

To that end, consider again the claim by McDowell I quoted earlier which was supposed to make intelligible why it seems incoherent to ascribe to a subject knowledge of the fact that p if all we could assume about this subject is that he possesses inconclusive warrants in support of believing p. Thus McDowell claimed that the notion of knowledge based on one's having inconclusive warrants would

entail that this knowledge “could be constituted by being in a position, which for all one knows, [is compatible with the falsity of p]. And that seems straightforwardly incoherent”. In our discussion of the first sense of one’s having conclusive perceptual warrants, I conceded that that claim is at least *prima facie* quite plausible. However, it also seems to me that this claim is plausible at all only on the assumption that one’s knowing that p is based in some way of one’s having knowledge of an independent fact.

The problem now is that the conception of knowledge based on the kind of conclusive warrants – where this notion is understood in the second of two senses – that McDowell thinks perceptual experiences can provide is not a kind of knowledge that is based on knowledge of an independent fact. For it should be clear here that when one forms the belief that p on the basis of one’s being in a state in which the obtaining of the fact that p becomes transparent to one, one does not form that belief on the basis of one’s prior knowledge of the fact that p; this, of course, would not give us a satisfying picture of how knowledge comes about since it already presupposes one’s having the relevant piece of knowledge. Hence, the assumption has to be that when one forms the belief that p on the basis of one’s being in a conceptual-experiential state in which the obtaining of the fact that p becomes transparent to one that this knowledge is based directly on the conceptual-experiential state. But this also means that the question of whether one’s knowledge is a matter of one’s being in a position that for all one knows is incompatible with the fact the fact that p simply does not arise.

The same point essentially applies to what we might consider to be an alternative to McDowell’s. Thus suppose that what supports one’s perceptual knowledge of an external world proposition p is not that one has conclusive warrant in the form of

one's being in a state in which the obtaining of the fact that p is transparent to one but rather that one has some justification for believing p in the form of one's being in a state that perceptually e-represents the obtaining of the fact that p. The idea is that, in the right sort of circumstances, one's forming the belief that p in this way can constitute one's perceptually knowing that p. Crucially, this is not a situation in which one forms the belief that p on the basis of one's having knowledge of some independent fact. Thus, there does not seem to be any relevant sense in which one is in a position that for all one knows is incompatible with the falsity of p. As such, we can avoid McDowell's charge that ascribing knowledge in this case is incoherent.

Consider now the other argument which McDowell earlier directed at the hybrid view. Recall that on this view, one's knowledge of an environmental proposition p was supposed to be constituted by one's having justification for believing (a) that it perceptually seems to one as if p and (b) that such seemings are a reliable indication that p is the case. McDowell's objection to this view was that it could not coherently conceive of knowledge as a matter of a standing in the space of reasons at all. For one implication of this view, as McDowell pointed out, was that two subjects could be alike in their (purported) standing in the space of reason and yet only of them know the fact that p. This, in turn, was supposed to entail that one's believing that p on the basis of one's having justification for believing (a) and (b) could not be a matter of a standing in the space of reasons at all.

Suppose then that we apply a similar reasoning to the view that one's knowledge of an environmental fact that p can be a matter, qua having justification for believing p, of one's being in a state of perceptually e-representing the obtaining of the fact that p. Again, it would be possible on this view for two subjects to be alike in terms of the justification they possess for believing p and yet for only one of them

to have knowledge of p. And so, McDowell might conclude once more that one's knowledge of p cannot be a matter of a standing in the space of reason at all. Since, however, knowledge is essentially a standing in the space of reasons, that simply entails that on the view under consideration knowledge is not possible at all.

In responding to this argument, we should note first of all that, in a sense, McDowell's conclusion obtains quite independently anyway. Thus, as I explained in the previous chapter, the state of perceptually e-representing the obtaining of the fact that p is not a state in which one has a reason for believing p, if the notion of a reason is one of an actual fact that one can perceive. Hence it follows that on the current view of perceptual knowledge such knowledge cannot indeed be a standing of a space of reasons if such a standing requires the possession of a reason *qua* perceived fact.

What does not follow, however, is that on this view perceptual knowledge cannot reasonably be considered a standing in a space of reasons in any sense at all. For first of all, as I suggested in the previous chapter, although the state of perceptually e-representing that p is not a state in which one possesses a reason *qua* experienced fact, it is a state that can nonetheless function as a reason² in Ginsborg's sense. Second, I would also suggest that the state of perceptually e-representing that p, in the absence of reasons for thinking that one is say, hallucinating, essentially puts one in a position to be able to properly justify one's believing p on the basis of claiming that one sees that p. In other words, such a state can put one in a position to have a reason¹ in Ginsborg's sense. From the fact that perceptual e-representations can amount to reasons² for beliefs and the fact that they can put one in a position to have a reason¹ for believing p, it follows that one's believing that p on the basis of one's

having a relevant perceptual e-representation can be a matter of a standing in the space of reasons in some sense or another.

McDowell could, of course, attempt to undermine this contention by appealing once again to the thought experiment involving two subjects who are alike in their (purported) standings in the space of reasons but of whom only one has knowledge. But as before, this thought experiment would show that perceptual knowledge cannot be a standing in the space of reasons when it is based on one's having a perceptual e-representation only if one already accepted that perceptual knowledge requires conclusive reasons in McDowell's sense. McDowell, however, simply does not offer an independent argument for this conclusion.

Chapter 4

Introduction

We have been looking at various possible explanations for what I have called the Basic Principle about Perceptual Justification. To recall, this principle, in its doxastic form, states that:

If a subject S has a perceptual experience as of a mind-independent object an x being F (or in which it appears to him as if an x is F), and forms the belief that an x is F on the basis of having an experience of this phenomenological sort, then (perhaps provided certain further conditions obtain) S's belief that an x is F is *prima facie* justified for S.

The concern so far has been with whether this principle could be explained by saying that it is the perceptual experience itself that contributes to the epistemic justification of the subject's belief and whether it does in virtue of the phenomenal character that it has. The suggestion, developed in chapter 2, was that we should think of the subject's belief as being justified on the basis of his perceptually e-representing the obtaining of the fact that an x is F (or the fact that makes true the proposition that an x is F), where this latter state is constituted by the subject's having a perceptual experience of the relevant phenomenological sort.

However, for all that we have said, the claim that the Basic Principle is true as a matter of the subject's perceptually e-representing the obtaining of the fact that

an *x* is *F* leaves open the possibility that further conditions have to be met in order for the subject to have epistemic justification for his belief that an *x* is *F*. In particular, the question remains whether the subject can be justified in his perceptual belief on the basis of having the perceptual experience that he does without having independent justification for further beliefs. In brief, the question remains whether the Basic Principle can be true as a matter of the subject's possessing immediate perceptual justification for his belief.

But why should we think at all that when a subject has justification for his belief that an *x* is *F* on the basis of perceptually e-representing the obtaining of the fact that an *x* is *F*, that his having this justification presupposes that he has independent justification for further beliefs? As we saw in the previous chapter, we might be persuaded by elements in the sceptic's line of reasoning into thinking that one can have a justified perceptual belief only to the extent that one has some a priori justification for believing that one is not a brain-in-a-vat. However, there is another line of argument that has been the concern of recent discussions in epistemology which might lead us into thinking that one cannot have a justified perceptual belief unless one also has some independent justification for believing that one's perceptual experiences are generally reliable. It is this line of argument that I will be addressing in this chapter.

The issues here have been brought into focus in a paper by Stewart Cohen.²⁷ Cohen presents these issues in terms of knowledge rather than justification, and I will follow him in this. We can begin with the plausible thought that we can come to know many things on the basis of perceptual experiences without first knowing that our perceptual experiences are generally reliable. A young child, for example, may

²⁷ (Cohen, *Basic Knowledge and the Problem of Easy Knowledge*, 2002)

know on the basis of looking at a table that the table is red without having prior knowledge, for which he may lack the requisite conceptual resources, that his visual experiences are a reliable source of information. Following Cohen, I will call all knowledge that is based on perception without prior knowledge that perceptual experiences are reliable basic (perceptual) knowledge. Any theory of perceptual knowledge that accepts the possibility of basic knowledge thus rejects the following principle:

PR If perceptual experiences provide S with knowledge that p, then S knows (antecedently) that perceptual experiences are reliable.

Thus according to what we might call a *foundationalism* (about perceptual knowledge), a subject can come to know that p on the basis of some grounds g, which provides him with justification for believing p, without prior knowledge that g is a reliable indication that p is the case.²⁸ For example, on this view, a subject can come to know that the table is red on the basis of a perceptual experience in which it appears to him as if a table is red – which provides him with justification for believing that a table is red – without having prior, independent knowledge that a perceptual experience in which it appears to him as if a table is red is a reliable indication that it is red. Note that the idea that the perceptual experience is the justificatory basis for his knowledge that a table is red need not mean that S forms the belief that a table is red through an inference from the belief that he is having a perceptual experience in which it appears to him as if a table is red. Rather, the

28 Ibid, 310.

perceptual experience can provide S with immediate, non-inferential knowledge that the table is red. As I understand the view, the perceptual foundationalist holds that the perceptual experience can provide the subject with immediate justification for believing that a table is red which, in the right circumstances, is sufficient, qua justification, for him to know that a table is red without his having prior knowledge that having a perceptual experience in which it appears to him as if a table is red is a reliable indication that a table is red. The perceptual foundationalist therefore rejects the following version of the PR principle:

PR1 If S knows that an x is F on the (justificatory) basis of having a perceptual experience in which it appears to him as if an x is F, then S has prior knowledge that his having a perceptual experience in which it appears to him as if an x is F is a reliable indication that an x is F.

The problem with any view that allows for basic knowledge, so Cohen argues, is that it also allows for the possibility of 'easy knowledge'. Again following Cohen, I will refer to this problem as the 'Problem of Easy Knowledge'. Cohen identifies two such problems. The first of these supposedly derives from the acceptance of both the possibility of basic knowledge and the following closure principle.

Closure If S knows that p, competently deduces q from p, and thereby comes to believe q while retaining knowledge of p, then S knows q.²⁹

²⁹ This is how the Closure Principle is defined by John Hawthorne. See (Hawthorne, 2004, p. 34). This definition is different from the definition Cohen employs in his paper.

Suppose then that S has basic visual knowledge that the table is red. In particular, S knows that the table is red simply on the basis of having the relevant visual experience. Given closure, S should then be able to deduce and thereby come to know, for example, that the table is not white cleverly illuminated to look red. Thus, having as justificatory grounds merely a perceptual experience in which it appears to him as if a table is red, S can come to know that the table is red and subsequently that the table is not white cleverly illuminated to look red. This form of 'easy knowledge', however, seems to be entirely unacceptable. On the assumption that the closure principle is undeniable, and that basic knowledge entails the possibility of easy knowledge, it would seem to follow that S does not have basic visual knowledge of the fact that the table is red after all. Furthermore, since the proposition that a subject can have immediate perceptual justification for believing a proposition p entails that he can have basic knowledge that p and therefore easy knowledge of a relevant proposition q, if easy knowledge is not possible then neither is immediate perceptual justification.

The second problem is that basic knowledge seems to allow for a process called 'bootstrapping'.³⁰ Thus suppose again that at a certain time t, S knows that the table is red on the basis merely of a perceptual experience in which it appears to him as if a table is red. Certainly, S can also come to know based on a bit of introspection that he is having this kind of perceptual experience. So at t, S can know that the table is red and that he is having a perceptual experience in which it appears to him as if a table is red and conclude on this occasion that the perceptual experience accurately indicates that the table is red. Continuing this process on subsequent occasions, S would seem to be able to gather more and more evidence which would ultimately

30 See also (Fumerton, 1995) and (Vogel, Reliabilism Levelled, 2000)

provide him with sufficient justification for coming to know that his having a perceptual experience in which it appears to him as if a table is red is a reliable indication that the table is red. Again, this seems an unacceptable result. On the assumption that inductive knowledge is generally possible and that there is nothing problematic with S's coming to know on suitable occasions that he is having a perceptual experience in which it appears to him as if a table is red, it would seem to follow that S cannot have basic visual knowledge of the fact that the table is red to begin with. And once again, since the proposition that a subject can have immediate perceptual justification for believing a proposition *p* entails that he can have basic knowledge that *p* and therefore acquire knowledge, by bootstrapping, of the proposition that his perceptual experience is reliable, if knowledge by bootstrapping is not possible then neither is immediate perceptual justification.

At this point, we still need to determine how exactly basic knowledge leads to the easy knowledge problem in the two ways described and, relatedly, how acceptance of the PR principle can avoid the problem in the case of perceptual knowledge. For now we should note that there seem to be good arguments for thinking that the PR principle is untenable. First, as we have seen, acceptance of the PR principle would seem to make it impossible for us to explain how young children can have perceptual knowledge. But, perhaps more importantly, it would also seem to lead to skepticism about perceptual knowledge. Thus suppose that the principle is correct. The question then is how a subject could know that his perceptual experiences are generally reliable. It would seem that the only plausible way for him to have this knowledge is through some inference from propositions that he knows, and has justification for believing, on the basis of perceptual experiences. If that is

the case, however, then a subject would never attain any perceptual knowledge since he would already require some perceptual knowledge before first acquiring it.³¹

This line of reasoning may ultimately lead one to accept that easy knowledge, though seemingly problematic, is possible after all. That is, if, on the one hand, accepting the PR principle leads to skepticism whereas, on the other hand, denying the PR principle and allowing for basic knowledge leads to the possibility of easy knowledge, then denying the PR principle and allowing for easy knowledge may seem to be the lesser of two evils.³² Still, occupying the latter position may seem far from satisfying. It is worthwhile to consider, therefore, whether the rejection of the PR principle and the acceptance of basic knowledge really lead to the possibility of easy knowledge.

The Closure Problem

I first turn to the easy knowledge problem for perceptual foundationalism that derives from the closure principle. As we have seen, according to Cohen, it is the theory's rejection of the PR1 principle combined with acceptance of the closure principle that leads to the possibility of acquiring easy knowledge. Thus, if a subject S1 can come to know that the table is red on the basis merely having the relevant sort of perceptual experience, then given closure it seems that S1 can come to know that the table is not white cleverly illuminated to look red on the basis of a simple deduction without needing any further evidence in support of the deduced

31 See (Cohen, Basic Knowledge and the Problem of Easy Knowledge, 2002, p. 309).

32 See (Van Cleve, 2003)

proposition. First, let's ask whether accepting the PR1 principle can avoid the possibility of acquiring easy knowledge in this way. Of course, if accepting PR1 entails skepticism, then a subject would not know that the table is red to begin with and therefore would be unable to acquire easy knowledge of the proposition that the table is not white cleverly illuminated to look red. But suppose for the moment that PR1 does not entail skepticism. Then, on a given occasion, a subject S2 can know that the table is red on the basis having a perceptual experience in which it appears to him as if a table is red and his knowledge that this perceptual experience is a reliable indication that the table is red. Given the closure principle, he can then deduce and thereby come to know that the table is not white cleverly illuminated to look red.

Now is S2's way of acquiring knowledge of this fact any more acceptable than S1's way of acquiring knowledge of the same fact? One reason why we might think that S2's way of acquiring knowledge is legitimate whereas S1's way is not is that S2, unlike S1, is in possession of additional evidence from which he can infer that the table is not white cleverly illuminated to look red.

Thus S2, unlike S1, can reason as follows:

- (1) I have an experience in which it appears to me as if the table is red.
- (2) My having this kind of perceptual experience is reliable indication that the table is red.
- (3) If my having this kind of perceptual experience is a reliable indication that it is red, then the table's looking red to me is a reliable indication that it is not white cleverly illuminated to look red.
- (4) My having this perceptual experience is a reliable indication that the table is not white cleverly illuminated to look red.

(5) Therefore: the table is not white cleverly illuminated to look red.

Given our assumptions each of the premises can be known by S2. (1) simply expresses a piece of introspective knowledge. (2) is a piece of knowledge we assume S2 to have given our acceptance of the PR1 principle. (3) is a straightforward conceptual truth and (4) follows from (2) and (3). So S2 can know that the table is not white cleverly illuminated to look red without having to make an inference from the proposition that the table is red. This does not rule out of course that S can come to know that the table is not white cleverly illuminated to look red on the basis of a deduction from the proposition that the table is red, which he knows to be true. But this deduction will be supported by S's possession of evidence that provides justification for both the belief that the table is red as well as the belief that the table is not white cleverly illuminated to look red. And this fact may be enough to render the deduction a legitimate way of coming to know the conclusion. It would seem therefore that acceptance of PR1, even if it does not lead to skepticism, can solve the problem of easy knowledge.

The next question we need to address is whether the rejection of PR1 entails the possibility of easy knowledge. As we have seen, the foundationalist rejects PR1 by claiming that a subject can know that the table is red on the basis of having a perceptual experience in which it appears to him as if the table is red without having prior knowledge that having this sort of perceptual experience is a reliable indication that it is red. But, as we have also seen, if a subject can know that the table is red merely on the basis of the perceptual experience, then it seems that the subject can come to know that the table is not white cleverly illuminated to look red in an all too easy way. However, it is important to note that a rejection of PR1 does not entail that

a subject can come to know that the table is red *merely* on the basis of the perceptual experience. One may hold, for example, that in order for a subject to know that the table is red on the basis of a perceptual experience, he also needs to be in possession of additional evidence where this further evidence does not consist in his knowing that his having the perceptual experience is a reliable indication that the table is red. If so, a theory that insists on the subject's possessing this further evidence may potentially avoid the problem of easy knowledge as it arises from the closure principle without at the same time endorsing the PR1 principle.

Suppose then that we accept the following condition on S's knowing that the table is red:

IND If S knows that a table he is looking at is red merely on the basis of having a perceptual experience in which it appears to him as if a table is red, S needs to know independently of his knowing that the table is red on the basis of having the perceptual experience that the table is not white but cleverly illuminated to look red.

[I take it that IND is an instance of the more general principle that: If S knows that an x is F merely on the basis of having a perceptual experience in which it appears to him as if an x is F, then S needs to know independently of his knowing that an x is F on the basis of that perceptual experience that he the situation he is experiencing is not a situation in which no x is F but which is subjectively indistinguishable for him from the situation in which it perceptually appears to him as if an x is F and an x is F.]

Now since a subject's knowledge that the table is red presupposes his having independent knowledge that the table is not white but cleverly illuminated to look red, he is then obviously no longer in a position to acquire knowledge of the fact that the table is not white but cleverly illuminated to look red. So in this case, easy knowledge will not be available to him. Crucially, however, the IND principle is neutral with respect to the PR1 principle. As stated the IND principle does not entail that S's independent knowledge that the table is not white but cleverly illuminated to look red, which he needs in order to know that the table is red on the basis of having the relevant perceptual experience, has to be supported by knowledge that having that perceptual experience is a reliable indication that the table is red or by his knowledge that having that perceptual experience is a reliable indication that it is not white cleverly illuminated to look red. And since the IND principle is consistent with the rejection of the PR1 principle, it is also consistent with the possibility of basic knowledge. Therefore, it can potentially explain how the possibility of basic knowledge need not entail the possibility of easy knowledge.

However, the IND principle is too strong to be plausible. Once again, young children prove a case in point. Still, we may find similar principles that are more restrictive and specifically relevant to the cases at hand. In this regard, it is important to recall that, in accordance with the closure principle, for a subject to come to know that the table is not white cleverly illuminated to look red by deducing this proposition from the proposition that the table is red, which he knows to be true, he needs to retain his knowledge that the table is red throughout the process. Now there are of course many different ways in which a subject, once having acquired knowledge that the table is red merely on the basis of it's looking red to him, can lose

this piece of knowledge. For example, he may hear testimony from a reliable source that there are numerous tables in his current environment that are white but cleverly illuminated to look red. Under such a circumstance, his perceptual experience can no longer function as adequate justification for sustaining S's knowledge that the table is red. I take it that this is because S is now in possession of evidence for the claim that having the perceptual experience, in his current environment, is not a reliable indication that the table is red or that the table he is looking at may indeed be white cleverly illuminated to look red, evidence which undermines or defeats whatever epistemic support the perceptual experience might otherwise provide. In this case, if S is to retain his knowledge that the table is red, he needs some independent grounds for dismissing the hypothesis that the table is white cleverly illuminated to look red. But given that his retaining the knowledge that the table is red now depends on his having independent grounds for dismissing the hypothesis that the table is white cleverly illuminated to look red, it is no longer possible for him to acquire easy knowledge of the fact that the table is not white cleverly illuminated to look red. Thus it seems plausible to accept the following principle:

IND1 If S knows that the table is red merely on the basis of having a perceptual experience in which it appears to him as if a table is red and subsequently acquires evidence in support of the claim that the table is white cleverly illuminated to look red, then S can retain his knowledge that the table is red only if S has some independent grounds for knowing that the table is not white but cleverly illuminated to look red.

[Generally: If S knows that an x is F merely on the basis having a perceptual experience in which it appears to him as if an x is F and subsequently acquires evidence in support of the claim that the situation he is experiencing is a situation Sn in which no x is F but which is phenomenologically indistinguishable for him from the situation in which he is having this perceptual experience and an x is F, then S can retain his knowledge that an x is F only if S has some independent grounds for knowing that the situation he is experiencing is not a situation Sn]

Now, crucially, it also may seem plausible to some to think that S's knowledge can be defeated in a similar fashion even if S is not provided with positive evidence for thinking that his visual experiences are not relevantly reliable. To illustrate, consider Cohen's example:

Suppose my son wants to buy a red table for his room. We go in the store and I say, 'That table is red. I'll buy it for you.' Having inherited his father's obsessive personality, he worries, 'Daddy, what if it's white with red lights shining on it?' I reply, 'Don't worry – you see, it looks red, so it is red, so it's not white but illuminated by red lights.'³³

In this case, the son does not provide Cohen with any positive evidence for believing that the perceptual experience they are having may not be a reliable indication that it is red. Nor does he provide Cohen with positive evidence for believing that the table they are looking at is indeed white but cleverly illuminated to look red. The son is

33 (Cohen, *Basic Knowledge and the Problem of Easy Knowledge*, 2002, p. 314).

simply asking whether the table may not be red but white cleverly illuminated to look red. Still, as Cohen suggests, it seems unacceptable to think that under the given circumstances, he could come to know that the table is red by reasoning the way he does. We might explain this by saying that although the son does not present Cohen with evidence for believing that the table they are looking at is actually white cleverly illuminated to look red, the son does make Cohen aware of a certain kind of error-possibility. That is, he raises the possibility that the table merely looks red to them when in fact it is not. Just like in the previous case, we might then think that in order for Cohen to retain his knowledge that the table is red, he needs to have some independent grounds for dismissing the hypothesis that the table is white cleverly illuminated to look red. Consequently, as in the previous case, given that in order for Cohen to retain his knowledge that the table is red, he needs to have independent grounds for dismissing the hypothesis that the table is white cleverly illuminated to look red, he is no longer in a position to come to know that this hypothesis does not obtain by deducing its falsity from the proposition that the table is red. Cohen's example therefore seems to make plausible the following principle:

IND 2 If S knows that the table is red merely on the basis of his having a perceptual experience in which it appears to him as if the table is red and subsequently entertains the possibility that the table is white but cleverly illuminated to look red, then S can retain his knowledge that the table is red only if S has some independent grounds for knowing that the table is not white cleverly illuminated to look red.

[Generally: If S knows that an x is F merely on the basis of his having a perceptual experience in which it appears to him as if an x is F and subsequently entertains the possibility that the situation he is experiencing is not a situation S_n in which no x is F but which is phenomenologically indistinguishable for him from the situation in which an x looks F to him and is F, then S can retain his knowledge that an x is F only if S has some independent grounds for knowing that the situation he is experiencing is not a situation S_n]

So we have two kinds of cases in which a subject starts off with having knowledge that the table is red merely on the basis of having a perceptual experience in which it appears to him as if the table is red but subsequently ends up being in no position to rely on this knowledge to come to know that the table is not white cleverly illuminated to look red. The question is whether the principles we have formulated provide us with a sufficient framework to solve the first of the two easy knowledge problems. If yes, the suggestion would have to be that a subject's acquisition of easy knowledge of the proposition that the table is not white cleverly illuminated to look red based on his knowledge that the table is red is not possible for this reason: while a subject can know that the table is red merely on the basis of having the relevant sort of perceptual experience, he cannot retain this knowledge in such a way that it could form the basis for him to come to know that the table is not white cleverly illuminated to look red. The reason for that in turn is that a subject who knows that the table is red on the basis of it's looking red to him cannot become aware that the proposition that the table is red entails the proposition that the table is not white cleverly illuminated to look red without thereby (a) acquiring evidence that

undermines the epistemic justification the perceptual experience may otherwise give him for believing that the table is red or (b) entertaining the possibility that the table merely looks red when actually it is not.

Now it seems to me that the two principles at least capture those cases that we could imagine to somewhat plausibly occur in real people's lives. That is, it would seem to be extremely odd for people who believe that a table they are looking at is red when the table looks red to them to even entertain the proposition that the table is not white cleverly illuminated to look red unless they were already somehow concerned with whether or not the table might merely appear to be red when in fact it wasn't. The only problematic case would seem to be one in which a subject who knows that a table he is looking at is red merely on the basis of having a perceptual in which it appears to him as if the table is red for no particular reason at all deduces and thereby comes to believe that the table is not white cleverly illuminated to look red. Still, there is a sense in which it seems to be epistemically illegitimate for him to come to believe that the table is not white cleverly illuminated to look red in this way if the only relevant justification he possesses is the relevant perceptual experience. After all, if he is in possession of no other relevant evidence, the perceptual experience cannot indicate to him whether the table either is red and therefore not white cleverly illuminated to look red or white but cleverly illuminated to look red. This would mean that his evidential situation cannot epistemically support his belief that the table is not white cleverly illuminated to look red, which means in turn that he lacks knowledge of this proposition. Given the closure principle, however, his lack of knowledge that the table is not white cleverly illuminated to look red entails his lack of knowledge that the table is red. Thus, in order for him to retain his knowledge that the table is red, the subject needs to have some independent grounds

for knowing that the table is not white cleverly illuminated to look red. We can then propose the following principle:

IND3 If S knows that the table is red merely on the basis of the table's looking red to him and subsequently comes to believe that the table is not white cleverly illuminated to look red, then S can retain his knowledge that the table is red only if S has some independent grounds for knowing that the table is not red cleverly illuminated to look red.

[Generally: If S knows that an x is F merely on the basis of having a perceptual experience in which it appears to him as if an x is F and subsequently comes to believe that the situation he is experiencing is not a situation S_n in which no x is F but which is phenomenologically indistinguishable for him from the situation in which an x looks F to him and is F, then S can retain his knowledge that an x is F only if S has some independent grounds for knowing that the situation he is experiencing is not a situation S_n]

The appeal to the various IND principles, therefore, provides us with the means to rule out the possibility of easy knowledge that seemed to derive from the possibility of basic knowledge and the closure principle. The key point is that, in the relevant cases, a subject cannot retain his knowledge that the table is red without having independent grounds for knowing that the table is not white cleverly illuminated to look red. Depending on the particulars of the case, such independent grounds can

consist of background knowledge to the effect that it is very unlikely that the table one is looking at, though it looks red to one, is not red but cleverly illuminated to look red. For example, Cohen in the earlier example might point out to his son that it is very unlikely that furniture stores would cleverly disguise the real colors of their pieces. Crucially, however, such background knowledge need not amount to having knowledge that having a perceptual experience in which it appears to one as if the table is red in general is a reliable indication that the table is red. Thus, at least far as the first of the two easy knowledge problems is concerned, we can provide a solution to this problem without having to appeal to the problematic PR1 principle.

The Bootstrapping Problem

I now turn to the second easy knowledge problem, which is that basic knowledge seems to allow for bootstrapping. As before, I will address this problem with respect to perceptual foundationalism. To remind ourselves how this problem arises consider the following example. Suppose that, on a given occasion, a subject S knows that the table in front of him is red on the basis of having a perceptual experience in which it appears to him as if the table is red. Again, we assume that S has no prior knowledge that his having this perceptual experience is a reliable indication that it is red. Thus S's knowledge that the table is red is a piece of basic knowledge. It now seems that S is in a position to reason as follows:

- (1) The table is red.
- (2) I have a perceptual experience in which it appears to me as if the table is red.
- (3) I have this perceptual experience and the table is red.

(4) My having this perceptual experience accurately indicates that the table is red.

Apart from having basic knowledge of (1), S should know or easily be capable of knowing each of the other premises. Certainly, S can straightforwardly know (2) on the basis of introspection. (3) logically follows from (1) and (2) while (4) logically follows from (3). Hence S should be able to come to know (3) and (4) on the basis of a straightforward deduction from previously known premises.

Suppose then that S repeats this kind of reasoning on many other occasions in which he has basic knowledge that the table in front of him is red on the basis of having a perceptual experience in which it appears to him as if the table is red. In doing so, he ultimately seems to acquire evidence which is sufficient, qua evidence, for him to come to know that having this sort of perceptual experience is a reliable indication that the table is red. Thus, from this knowledge of (4) and his knowledge of many other propositions relevantly like (4)³⁴, which he can arrive at on the basis of the same kind of reasoning just outlined, he can inductively infer and come to know that

(5) My having a perceptual experience in which it appears to me as if a table is red is a reliable indication that the table is red.

On the face of it, this kind of bootstrapping process seems to be an illegitimate way of coming to know (5). The inference from (4),... to (5), however, is inductively valid. Thus, if S does not know (5), then it would seem that this must be because he

³⁴ I shall symbolise this sentence in this way from now on: '(4),...'

does not know (4),... after all. As we have seen, however, (4) logically follows from (3), so that if S does not know (4), he does not know (3) either. We have also seen that (3) logically follows from (1) and (2) and since, presumably, there is nothing problematic about S's introspective knowledge of (2), it follows from S's lack of knowledge of (3) that he lacks knowledge, or at least basic knowledge, of (1). Reflection on bootstrapping, therefore, seems to initiate what Jonathan Vogel calls a 'rollback' of S's presumed original basic knowledge of (1).

We can initially approach this second easy knowledge problem for perceptual foundationalism in the same way that we approached the first easy knowledge problem. Thus, we should ask, first, whether acceptance of PR1 – that is, acceptance of the principle that if S knows that the table is red on the basis of a perceptual experience in which it appears to him as if the table is red, he has to have prior, independent knowledge that his having this sort of perceptual experience is a reliable indication that the table is red – can avoid the bootstrapping problem and, second, whether a rejection of PR1 entails the possibility of bootstrapping. The answer to the first question is a straightforward yes. If S knows that the table is red on the basis of the right sort of experience only if he has prior, independent knowledge that his having this sort of perceptual experience is a reliable indication that it is red, then obviously he will no longer be able to acquire this piece of reliability knowledge by reasoning the way he does.

What then about the second question? Earlier we noted with respect to the first easy knowledge problem that the rejection of PR1, and therefore the acceptance of the possibility of basic knowledge, do not entail that S can know that the table is red merely on the basis of his having the relevant perceptual experience. Thus, it may be in some cases that although a subject S can have basic knowledge that the

table is red on the basis a perceptual experience, S also needs to be in possession of additional evidence distinct from the justification he has as a matter of having the perceptual experience but not consisting of some prior, independent knowledge that his having this sort of perceptual experience is a reliable indication that the table is red. In particular we saw that in the relevant cases, a subject can retain his basic perceptual knowledge that the table is red only if he also has independent evidence or grounds for knowing that it is not, for example, cleverly illuminated to look red. In this way, the requirement, in the relevant cases, that for S to retain his knowledge that the table is red he be in possession of such independent evidence or grounds for knowing the table is not red but cleverly illuminated to look red, provided some means for avoiding the first easy knowledge problem without having to appeal to the problematic PR1 principle.

Unfortunately, however, this kind of approach does not seem to be available when it comes to the second problem of easy knowledge. Thus suppose, for argument's sake, that the relevant cases can be restricted to those in which S actually engages in the kind of bootstrapping outlined earlier. What we want to say then, in line with our approach to the first easy knowledge problem, is that for S to have or retain knowledge of the premises, S has to be in possession of evidence or grounds in addition to the justification he has as a result of his having a perceptual experience in which it appears to him as if the table is red which would help explain either why S cannot come to know, by bootstrapping, that his having this perceptual experience is a reliable indication that the table is red or why the reasoning he performs is not an illegitimate way of coming to know the conclusion after all. The problem, however, is that it is difficult to see what such relevant additional evidence, if it does not consist of some prior, independent knowledge that his having a perceptual

experience in which it appears to him as if the table is red is a reliable indication that the table is red, could amount to. For if the relevant additional evidence does not consist of prior, independent knowledge of the conclusion of the bootstrapping reasoning, it obviously has to consist of evidence in support of one or several of the premises. Here the relevant premises are really (1) and (3) (or (4)). Of course, it follows that if S did have evidence in support of (1) or (3) independently of the justification he gains from looking at the table, then the reasoning from (1) – (4) to the conclusion (5) may be entirely appropriate. In that case, S would be accumulating evidence in support of the conclusion that his having a perceptual in which it appears to him as if the table is red a reliable indication that the table is red independently of the justification provided to him his perceptual experience. But it seems clear that such independent evidence cannot be had. For surely, such evidence would have to be a piece of empirical evidence. However, whatever empirical evidence S could gather in support of the belief that the table is red without relying on how things appear to him– for example, someone's testimony – that distinct piece of evidence in turn will ultimately have to be supported by evidence gained on the basis of someone or other looking at the table. So the requirement that S be in possession of additional evidence distinct from the justification afforded to him by his perceptual experience but not consisting of prior, independent knowledge that his having the perceptual experience is a reliable indication that it is red seems to get us nowhere.

From this we might conclude that, really, the only way of avoiding the second easy knowledge problem is to accept principle PR1. Hence upon reflection on the issue of bootstrapping, we do seem to be faced with the rather uncomfortable choice of having to accept either the possibility of easy knowledge or a problematic

epistemic principle that seems to lead to skepticism. I do think that this dilemma is real but before turning to how we should proceed from here, I want to address a somewhat different approach to the bootstrapping problem.

In a recent paper, Jonathan Vogel³⁵ has argued that the available options in the face of the second easy knowledge problem are not in fact exhausted by accepting the possibility of bootstrapping on the one hand and a problematic principle like PR1 on the other. Instead he contends that a subject can have basic knowledge of (1) and thereby be in a position to know the remaining premises without also thereby being in a position to come to know (5) by reasoning from (4),... to (5). Now it should be clear that what is problematic about the reasoning we are concerned with is that it involves a certain kind of epistemic circularity. Thus in forming the belief that his having a perceptual experience in which it appears to him as if the table is red is a reliable indication that the table is red, our subject S is in fact relying on his perceptual experience in support of the premises in his reasoning. However, Vogel suggests that a sound principle governing justification (and, by implication, knowledge) is something like the following:

NEC³⁶ We cannot acquire knowledge that our perceptual experience
 E is reliable by relying in part on E itself.³⁷

This principle certainly seems wholly plausible, and it is what seems to underlie the intuition that bootstrapping is an epistemically illegitimate procedure. But even if a

³⁵ (Vogel, Epistemic Bootstrapping, 2008)

³⁶ NEC = No Epistemic Circularity

³⁷ The principles that Vogel himself formulates and appeals to are these. (NRC) A belief that an epistemic rule R is reliable cannot be justified by the application of R. That is, neither the conclusion itself nor any belief which supports the conclusion may be justified in virtue of the application of R and (NSS) One cannot obtain...a justified or warranted believe that a belief source S is trustworthy by relying even in part on source S. (Ibid, p. 531)

principle like NEC is accepted, the question of course is how it can offer a viable option for addressing the second easy knowledge problem which avoids having to accept either the possibility of bootstrapping or a problematic principle like PR1. As Vogel sees it, the issue turns on whether the appeal to NEC can provide us with the resources to resist the rollback problem. Thus, as we have seen earlier, it seems entirely plausible to say that if S does not know (5), then this must be because he does not know (4). We have also seen, however, that if (4) is not known, then neither is (1), at least not in a basic way. Vogel's proposed solution to this problem is to say, first of all, that other things being equal, S's knowledge of (4), supported as it is by his basic knowledge of (1), does provide S with a warrant for believing, and thereby knowing, conclusion (5). A principle like NEC, however, can operate as a potential defeater to the justification S's knowledge of (4),...may otherwise provide for believing (5). For example, upon realizing that the reasoning he employs violates NEC, S would then no longer be in a position to come to know (5) by inferring (5) from (4),.... Crucially, however, this does not mean, that S thereby loses his knowledge of (4). The rollback problem, therefore, is resisted.

To illustrate his point, Vogel provides this analogy. Suppose that S knows that Tweety is a bird. This knowledge may well give him warrant for believing that Tweety flies. However, it fails to do so if S also knows that Tweety is a penguin. Thus his knowledge that Tweety is a penguin defeats the warrant that S, in knowing that Tweety is a bird, would otherwise have for believing that Tweety flies. Having his warrant for believing that Tweety flies defeated in this way, however, does not entail that S loses his knowledge that Tweety is a bird.³⁸

³⁸ Ibid, p. 535, n. 42

Now is this a good solution to the second easy knowledge problem? The most obvious problem with this solution is that it does not completely avoid the dilemma we were faced with earlier. This of course contradicts Vogel's advertised goal since his contention was that the options in the face of the second easy knowledge problem were not exhausted by accepting the possibility of bootstrapping on the one hand and a problematic principle like PR1 on the other. As it is now clear, however, given his own solution to the problem, Vogel has to accept that bootstrapping is possible in some cases, namely in those in which a principle like NEC does not operate as a defeater.

Still, despite this, Vogel also thinks that there are some cases in which a subject does have basic knowledge of (1) as well as knowledge of (2) – (4) but is not in a position to acquire the relevant piece of reliability knowledge by employing the bootstrapping reasoning. This means that at least in those cases, we can explain the impossibility of bootstrapping without appealing to a principle like PR1 but by appealing instead to a principle like NEC.

The first point to note here is that Vogel does not provide a full account of the circumstances in which a principle like NEC does operate as a defeater to the warrant S's knowledge of (4) may otherwise provide for believing (5). One type of case I mentioned earlier was a case in which S comes to realize that his reasoning violates NEC, and this seems to be the type of case that Vogel has in mind. So, focusing on this type of case, does it really make sense to say, as Vogel suggests, that S retains his justification for, and knowledge of, (4),... upon realizing that it cannot epistemically support belief in (5)? In other words, is this case sufficiently analogous to is the Tweety case?

There are two relevant comments Vogel makes in regard to this question. First, he says that: "...the specific flaw in the justification for [5] does not attach to [4]". (Vogel, *Epistemic Bootstrapping*, 2008, p. 536) In other words, while the (supposed) justification for (5) that the reasoning from (4),...to (5) provides is epistemically circular in that it violates NEC, the justification or warrant for believing (4) is not. This point is entirely straightforward: NEC is a principle concerning knowledge of reliability and thus, since (4) does not express a reliability claim, S's knowledge of or warrant for (4) cannot be subject to NEC. The implication of this point, however, or so Vogel suggests, is that while S may fail to have justification for believing (5) as a result of a violation of NEC, he need not thereby fail to have justification for believing (4). Hence, rollback is averted. Second, Vogel says: "would your discovery of bootstrapping...indicate that your original beliefs...were false or unsupported? It seems not. Why should a subsequent, untoward inference proceeding from those original beliefs affect their epistemic standing? More broadly, if [epistemic circularity] or anything else blocks the justification of a reliability belief downstream, that seems irrelevant to the justification of one's beliefs upstream. If this assessment is correct, then [epistemic circularity] can explain the unacceptability of bootstrapping without bringing on the rollback problem." (Ibid, 536) As I understand it, Vogel's second comment directly relates to the first comment. His point seems to be that if S comes to realize that the inference from (4),... to (5) is in violation of NEC and thereby acquires a reason for thinking that the inference is illegitimate, which is sufficient for the justification for believing (5) to be defeated, S need not thereby acquire a reason for thinking that there is no justification for believing (4). Presumably this is because, as the first comment makes clear, the belief in (4), unlike the belief in (5), does not violate

NEC. So there is nothing in S's realization that the bootstrapping reasoning violates NEC that would provide him with a reason for thinking that he has, after all, no justification for (4).

I believe that we can raise an objection to each of Vogel's two comments. One concern with the first comment is that while S's justification or warrant for (4) is not epistemically circular in the sense that it does not violate principle NEC, it does exhibit a different but similar, and seemingly equally crippling form of, epistemic circularity. Recall that (4) expresses S's belief that his having a perceptual experience in which it appears to him as the table is red [on the given occasion] is an accurate indication that the table is red. As we have seen this belief was based ultimately on the given perceptual experience. But certainly, if NEC is a plausible principle governing epistemic justification or warrant, then so is the following:

NEC1 We cannot acquire knowledge that a perceptual experience E accurately indicates the truth of p on a given occasion by relying on E itself.

Thus this principle states that S cannot come to know that his having a perceptual experience in which it appears to him as if the table is red is a correct indication that the table is red on a given occasion by relying in one way or another on that perceptual experience. It seems entirely plausible to assume that if NEC can operate as a defeater to the justification S's knowledge of (4) may provide for believing (5), NEC1 can operate as a defeater to the justification that S's knowledge of (1) may provide for believing (4) [via (2) and (3)]. Still, Vogel might say that NEC and NEC1 express different principles and as such, it is possible that S becomes aware of

the former without becoming aware of the latter. In such a situation S's justification for (5) would be defeated while his justification or warrant for (4) may well remain intact. The problem with this response, however, is that it is simply difficult to imagine a subject who is acute to the fact that his reasoning violates NEC without at the same time being acute to the fact that his reasoning also violates NEC1. Even if it is possible, in principle, that such a situation obtains what we could say, at the very least, is that if S really is in such a situation, he shows a significant epistemic failing with respect to his belief of (4), so that there is a real question as to whether or not this belief is genuinely justified and thereby amounts to knowledge.

This brings us to the second of Vogel's two comments. To repeat, Vogel's point, as I understood it, was that when S becomes aware of a violation of NEC in his reasoning and thereby acquires a reason for thinking that the reasoning is illegitimate, S need not thereby acquire a reason for thinking that there is no justification for believing (4) either. Hence while S's justification for (5) may be defeated, his justification for (4) may well remain intact. Now I have just given a reason for why Vogel's contention here is doubtful: for, as I suggested, it is difficult to see how S can be acute to a violation of NEC in his reasoning while not being acute to the violation of NEC1. So, contra Vogel, a defeat in S's warrant for (5) should also involve a defeat in his warrant for (4).

But there is another reason for doubting Vogel's contention. What we have so far is that when S becomes aware of the violation of NEC in his reasoning, he essentially becomes aware of the fact that his reasoning cannot provide him with justification for believing (5). But given our assumptions – namely that S has basic knowledge of premise (1) in his reasoning – S also has no further reasons he can avail himself of in support of (5). Thus, S should find himself in a situation in which

the rational thing for him to do is at the very least to withhold belief in (5). Withholding belief in (5), of course, is tantamount to withholding belief in the proposition that one's perceptual evidence is reliable. But if it is rational to withhold belief in the proposition that one's perceptual evidence is reliable, surely this justifies withholding any beliefs that are based on this evidence. Finally, it is plausible to think that if it is justified for S to withhold any beliefs based on his perceptual experience, then S is not justified in holding those beliefs, and this means that the original justification for those beliefs that S may have possessed is defeated.

Earlier Vogel asked whether 'your discovery of bootstrapping...[would]...indicate that your original beliefs...were false or unsupported?' There is a clear sense, contrary to Vogel's contention, that the answer to this question is yes. Given our assumptions, his discovery of bootstrapping leaves S with no justification for thinking that his perceptual experience is reliable, and that, if anything, should leave him with a sense that his original perceptual belief is not justified after all.

With this, we can conclude that Vogel has not managed to solve the second easy knowledge problem in a way that avoids the dilemma we were faced with earlier. For one, Vogel's solution entails that bootstrapping does lead to knowledge of reliability in some cases, so to that extent he simply embraces one of the horns of the dilemma. According to Vogel, it is only in some restricted cases of when S has basic knowledge of (1) and also has knowledge of (2)-(4), that S fails to be in position to acquire knowledge of (5) by reasoning the way he does. It is in those cases, that we can supposedly explain how reliability knowledge via bootstrapping is impossible, while retaining the assumption that the subject knows the premises of the bootstrapping reasoning, without at the same time appealing to a problematic

principle like PR1. But as I have tried to argue, Vogel's attempt to explain the impossibility of acquiring reliability knowledge via bootstrapping by appeal to a principle like NEC, while preserving the thought that the subject knows the premises of the bootstrapping reasoning, ultimately fails.

In the end, I think there are more general and fundamental reasons for thinking why Vogel's strategy is bound to fail. That is, while the appeal to a principle such as NEC may explain, or rather may make explicit, why we think bootstrapping is illegitimate, it cannot be used to explain how we can have basic knowledge while not being in a position to acquire reliability knowledge via bootstrapping. The reason for this is that anyone taking such an approach will have to accept that a subject can know the premises of an inductively valid argument but essentially still fail to know the conclusion. In other words, such an approach would be in contradiction with some inductive closure principle.

IC If S knows p, competently inductively infers q from p, and believes q on this basis while retaining knowledge of p, then S knows q.

Just like NEC, IC seems entirely plausible. But I suggest that it is the combination of these two principles that makes the second easy knowledge problem so seemingly intractable for us. For according to NEC, the bootstrapping reasoning cannot provide S with knowledge of reliability. But if the bootstrapping reasoning cannot provide S with knowledge of reliability, then given IC, the rollback problem seems to follow; that is, it seems to follow that S cannot know the premises of the bootstrapping reasoning either.

Knowing vs. Claiming to Know

The question that we raised earlier was whether in light of the bootstrapping problem there was any way to avoid the dilemma of having to accept either the possibility of easy knowledge or a principle like PR. The discussion of Vogel's attempt to solve the bootstrapping problem, however, shows the difficulty in trying to avoid this dilemma. Rather than exploring any further attempts, it might be worthwhile to determine, at this point, if it is possible at to accept one of the two horns of the dilemma in an intellectually satisfying way.

Now I take it here that acceptance of PR is more problematic than acceptance of the possibility of easy knowledge since the former, as we have seen, seems to imply the truth of scepticism about empirical knowledge. Thus if a choice has to be made between the possibility of easy knowledge on the one hand and the impossibility altogether of perceptual knowledge on the other, then the choice that we should make is clear enough. The only question remaining would then be whether acceptance of the possibility of easy knowledge could be made intellectually palatable. The task at hand here is to offer some considerations that might explain why, though entirely possible, the notion of easy knowledge by way of the bootstrapping process seems so counterintuitive.

To this end, consider once again principle NEC.

NEC We cannot acquire knowledge that our perceptual experience
 E is reliable by relying in part on E itself.

The suggestion earlier was that it is a principle of this kind that makes impossible knowledge of the reliability of perceptual experience by way of the bootstrapping process. Of course, if knowledge of the reliability of perceptual experience by way of the bootstrapping process is possible, then NEC cannot be a correct principle about perceptual knowledge and justification. However, to explain the counterintuitiveness of the possibility of easy knowledge and of the rejection of NEC, we can perhaps propose a somewhat different but related principle, concerning not knowledge but rather claims to knowledge.

NEC' We cannot appropriately claim to know that our perceptual experience E is reliable by relying in part on evidence gathered on the basis of E itself.

How should we understand this principle? To answer this, I want to turn some suggestions Duncan Pritchard makes in a different context about the conditions under which it is appropriate to claim to know a proposition p. Pritchard begins with this observation:

...we rarely convey our knowledge by making assertions which are prefixed with the phrase, 'I know'. Instead, one typically conveys one's knowledge of a proposition simply by asserting the proposition in question. Adding the further 'I know' phrase is rare, and standardly reflects not just emphasis but also an ability to resolve a particular challenge that has been raised. (Pritchard, 2008, pp. 302-3)

The suggestion here is, first of all, that assertions or claims to know a proposition *p* are typically directed at particular challenges to one's knowledge of *p*. Furthermore, Pritchard suggests, or seems to do so, that there are two key elements to such claims to knowing that *p*. First, in claiming to know that *p* in response to a specific challenge, one represents oneself as having some reasons for believing *p* that specifically address the particular challenge at hand. Second, in claiming to know that *p*, one represents oneself as having reasons for believing *p* that are stronger (or relevantly different) than the reasons one would originally need for the piece of knowledge itself.³⁹ These two elements to a claim to know a proposition *p* would seem to imply that such a claim is appropriate to the extent that one actually does have some reason for believing *p* that specifically address the particular challenge at hand and to the extent that one actually does have some reason for believing *p* that are stronger (or relevantly different) than the reason one would have originally needed for the piece of knowledge itself.

To illustrate these points, consider the following familiar example from the epistemological literature. Thus, suppose that while visiting the zoo with my son, I see a zebra in an enclosure marked 'Zebras'. I direct my son's eyes to the animal and say to him that that's a zebra. It is of course entirely plausible to think that my assertion here expresses my knowledge of the fact that the animal I am pointing to is in fact a zebra. Furthermore, we can assume that my ground for knowing this to be the case consists simply in my having a visual experience of the zebra. But suppose that my rather sceptical son challenges my assertion by asking whether the animal I am pointing to might not instead be a mule cleverly painted to look like a zebra. This

³⁹ Pritchard, p. 303

particular challenge by my son to my assertion that that's zebra thus consists of making me aware of a certain error possibility.

Pritchard's suggestion here is that in insisting that I know that that's a zebra in response to the challenge my son has raised, I will essentially represent myself as having some reason for believing that that's a zebra which speaks specifically to the error possibility being raised. As such, it would seem inappropriate in this case to claim that I know that that's a zebra if I could not defend this claim by saying anything other than that I can see that that's a zebra. Thus the grounds that I originally had for properly asserting, and indeed for knowing, that that's a zebra are in this case insufficient to support my response to my son's challenge. This is because in claiming to know that that's a zebra in response to that challenge, I represent myself as having some reasons for insisting that that's a zebra that specifically address the challenge at hand and therefore as knowing, for example, that the authorities at the zoo are unlikely to engage in such a conspiracy. Consequently, the grounds that make my claim to knowledge appropriate may be different than the grounds that supported my original piece of knowledge.

My suggestion is that we can apply a similar framework to assertions that one knows that a particular belief source such as one's perceptual experience is reliable. To illustrate this, consider the following case from Vogel's paper. Thus suppose that I am in the grocery store and perform a bootstrapping procedure that leads to my believing that my memory is reliable. Thus suppose that I remember that I am out of milk. We can assume that my memory can provide me with knowledge that I am out of milk independently of my already knowing that my memory is reliable. I also come to know through introspection that I remember that I am out of milk and conclude on the basis of this introspective knowledge and my memorial knowledge

that I am out of milk that my memory in this instance accurately indicates the fact that I am out of milk. Suppose then that I repeat this process with respect to other things that I correctly remember that I am out of and conclude and assert in the end that my memory is reliable. By hypothesis, my assertion expresses my knowledge that my memory is reliable. (Of course, there is a strong intuition here that this hypothesis is false.)

Next suppose that I am faced with a challenge to my assertion. We can assume that this challenge can simply consist of the question of how I know that my memory is reliable. Suppose, finally, that I respond to this challenge by insisting that I know that my memory is reliable and cite as a reason for believing that proposition to be true the bootstrapping reasoning I employed. Once again, there is a clear intuition here that it is inappropriate for me to insist that I know that my memory is reliable. We can explain this by saying that in responding to the challenge at hand by asserting that I know that my memory is reliable, I represent myself as having some evidence for thinking that my memory is reliable independently of evidence that I have gathered on the basis of that very belief source. Hence, intuitively, claims to know that a belief source is reliable seem to be governed by a general principle of which NEC' is a particular instance. This general principle would state:

NEC'' We cannot claim to know that a belief source S is reliable by relying on evidence acquired on the basis of S itself.

Recall now that we earlier accepted the hypothesis that I know on the basis of the bootstrapping reasoning that my memory is reliable. Given the assumption that my claim to know that my memory is reliable is inappropriate, this would be a case in

which by claiming to know that p I represent myself as having grounds in support of believing that my memory is reliable that are different than the grounds that would be required in simply coming to know that proposition to be true. Since I do not have such grounds, I am not in a position to appropriately assert that I know that my memory is reliable.

To sum up, we can say that knowledge of the reliability of perceptual experience by way of the bootstrapping process is counterintuitive because of the intuitiveness of a principle like NEC. The acceptance of the possibility of easy knowledge by way of bootstrapping, however, entails that NEC, though intuitive, is in fact false. The good news, though, is that we might be able to explain away that intuitiveness of NEC by appealing to a principle like NEC'. Thus we might be able to say that NEC is intuitive, even if false, because in judging NEC to be true we are implicitly making a correct judgment about a principle like NEC'. The consequence is that though knowledge of the reliability of perceptual experience by way of bootstrapping is strictly speaking possible, we cannot appropriately claim to have this knowledge if the only grounds we have for believing that our perceptual experiences are reliable is the bootstrapping reasoning itself.

Bibliography

- Alston, W. (1989). Foley's Theory of Epistemic Rationality. *Philosophy and Phenomenological Research* 50 (1) , 135-47.
- Bergman, M. (2006). *Justification Without Awareness*. Oxford: Oxford University Press.
- Bonjour, L. (1985). *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- Bonjour, L. (2004). In Search of Direct Realism. *Philosophy and Phenomenological Research* 69 (2) , 349-67.
- Brueckner, A. (2007). Content Externalism, Entitlement, and Reasons. In S. Goldberg, *Internalism and Externalism in Epistemology* (pp. 160-76). Oxford: Oxford University Press.
- Burge, T. (2003). Perceptual Entitlement. *Philosophy and Phenomenological Research* 67 (3) , 503-48.
- Cohen, S. (1984). Justification and Truth. *Philosophical Studies* , 279-295.
- Cohen, S. (2002). Basic Knowledge and the Problem of Easy Knowledge. *Philosophy and Phenomenological Research* 65 , 309-29.
- Davidson, D. (2000). A Coherence Theory of Truth and Knowledge. In S. Bernecker, & F. Dretske, *Knowledge - Readings in Contemporary Epistemology* (pp. 413-41). Oxford: Oxford University Press.
- Dretske, F. (1995). *Naturalizing the Mind*. Cambridge: MIT Press.
- Feldman, R. (1989). Foley's Subjective Foundationalism. *Philosophy and Phenomenological Research* 50 (1) , 149-58.
- Foley, R. (1987). *The Theory of Epistemic Rationality*. Cambridge, MA: Harvard University Press.
- Foley, R. (2003). Epistemic Rationality as Invulnerability to Self-Criticism. In S. Luper, *Essential Knowledge*. Longman.
- Fumerton, R. (1995). *Metaepistemology and Skepticism*. London: Rowman and Littlefield.
- Ginsborg, H. (2006). Reasons for Belief. *Philosophy and Phenomenological Research* 72 (2) , 286-318.
- Goldman, A. I. (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.

- Goldman, A. I. (1988). Strong and Weak Justification. *Philosophical Perspectives* 2 , 51-69.
- Harman, G. (1990). The Intrinsic Quality of Experience. *Philosophical Perspectives* 4 , 31-52.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford : Oxford University Press.
- Henderson, D., & Horgan, T. (2007). Some Ins and Outs of Transglobal Reliabilism . In S. C. Goldberg, *Internalism and Externalism in Semantics and Epistemology* (pp. 100-130). Oxford: Oxford University Press.
- Jackson, F. (1977). *Perception: A Representative Theory*. Cambridge: Cambridge University Press.
- Martin, M. (2002). The Transparency of Experience. *Mind and Language* 17 (4) , 376-425.
- Martin, M. (2004). The Limits of Self-Awareness. *Philosophical Studies* 120 , 37-89.
- Martin, M. (2006). On Being Alienated. In T. Gendler, & J. Hawthorne, *Perceptual Experience* (pp. 354-410). Oxford: Oxford University Press.
- McDowell, J. (1982). Criteria, Defeasibility, and Knowledge. *Proceedings of the British Academy* , 455-79.
- McDowell, J. (1994). *Mind and World*. Cambridge, MA: Harvard University Press
- McDowell, J. (1995). Knowledge and the Internal. *Philosophy and Phenomenological Research* 55 (4) , 877-93.
- McDowell, J. (2002). Knowledge and the Internal Revisited. *Philosophy and Phenomenological Research* 64 (1) , 97-105.
- McDowell, J. (2008). The Disjunctivist Conception of Experience as Material for a Transcendental Argument. In A. Haddock, & F. Macpherson, *Disjunctivism: Perception, Action, Knowledge* (pp. 376-89). Oxford: Oxford University Press.
- Noe, A. (n.d.). Thought and Experience.
- Peacocke, C. (1983). *Sense and Content*. Oxford: Oxford University Press.
- Plantinga, A. (1993). *Warrant and Proper Function*. Oxford: Oxford University Press.
- Pritchard, D. (2008). McDowellian Neo-Mooreanism. In A. a. Haddock, *Disjunctivism: Perception, Action, Knowledge* (pp. 283-310). Oxford: Oxford University Press.

- Pryor, J. (2004). What's Wrong with Moore's Argument? *Philosophical Issues* 14 , 349-378.
- Robinson, H. (1994). *Perception*. London: Routledge.
- Searle, J. (1983). *Intentionality*. Cambridge: Cambridge University Press.
- Snowdon, P. (1980). Perception, Vision, and Causation. *Proceedings of the Aristotelian Society* , 175-92.
- Sosa, E. (1991). *Knowledge in Perspective - Selected Essays in Epistemology*. Cambridge: Cambridge University Press.
- Strawson, P. (1979). Perception and Its Objects. In G. MacDonald, *Perception and Identity: Essays Presented to A.J. Ayer with His Replies* (pp. 41-60). London: Macmillan.
- Swain, M. (1989). On Richard Foley's Theory of Epistemic Rationality. *Philosophy and Phenomenological Research* , 159-68.
- Tye, M. (1995). *Ten Problems of Consciousness*. Cambridge : MIT Press.
- Van Cleve, J. (2003). Is Knowledge Easy or Impossible? Externalism as the Only Alternative to Skepticism. In S. Luper, *The Sceptics: Contemporary Essays* (pp. 45-59). Aldershot: Ashgate Publishing.
- Vogel, J. (2000). Reliabilism Levelled. *Journal of Philosophy* 97 , 602-27.
- Vogel, J. (2008). Epistemic Bootstrapping. *Journal of Philosophy* 105 (9) , 518-39.
- Wright, C. (2002). (Anti-) Sceptics Simple and Subtle: G.E. Moore and John McDowell. *Philosophy and Phenomenological Research* 65 (2) , 330-48.
- Wright, C. (2004). Warrant for Nothing (and Foundations for Free)? *Proceedings of the Aristotelian Society, Supplementary Volume* 78 , 167-212.
- Wright, C. (2008). Comments on John McDowell's 'The Disjunctive Conception of Experience as Material for a Transcendental Argument'. In *Disjunctivism: Perception, Action, Knowledge* (pp. 390-404). Oxford: Oxford University Press.