

[For John Greco (ed.) *The Oxford Handbook to Scepticism* (Oxford University Press, forthcoming.)

Disjunctivism and Scepticism

Alan Millar

1. Introduction

Disjunctivism is a way of thinking about perceptual experience. In one central form it is the view that judgements that characterise how it appears to a subject as if things are have disjunctive truth-conditions. To appreciate what is distinctive about disjunctivist thinking and why it is interesting, it is necessary to consider a *traditional approach in the theory of perception*. This will provide us with the background needed to assess the epistemological significance of disjunctivism and its bearing on philosophical scepticism.

2. Disjunctivism the traditional conception of experience

When I perceive some physical object, say, a cat on a chair, I have a visual experience such that it looks to me just as if a cat is on a chair before me. Philosophers have commonly thought that the very experience I have as I look at the cat could have been had by me even if no cat were there. Assuming that seeing is a 'success' notion, so that I do not see an F unless that F is there, it follows that merely having the experience cannot be sufficient for seeing the cat. But it also seems possible that in a circumstance in which the cat *is* there, I still might not see it despite having an appropriate experience. If I have the experience because my brain has been cleverly manipulated by neuro-physiologists, and the experience just happens to be such that it is as if a cat is there, then surely I do not see the cat. Under the conditions of the example there is nothing to connect my having the experience with the presence of the cat. Some connection is required for perception. A plausible response to this, and one characteristic of the traditional approach in the theory of perception, is to suppose that I see the cat only if,

- (a) I have an appropriate experience,
- (b) a cat is there, and further,
- (c) my having the experience is causally dependent on there being this cat before me.

This is the line taken by defenders of *the causal theory of perception*.¹

The ideas that I have just described form part of a traditional way of thinking about what it is to perceive an object. What I shall call *the traditional conception of experience* is the conception of experience that forms part of this way of thinking. Some terminology will help to sharpen up the conception. I shall call experiences *perceptual* when their subjects come to have them through the normal operation of the relevant sensory modality in an episode in which they perceive something. I shall say that experiences are *phenomenally visual* if they have the sort of phenomenal features that characterise visual perceptual experiences—the kind of features that perceivers can capture by describing how it looks as if things are. Experiences can be phenomenally visual by this account even if they are not perceptual. For traditionalists the very experience I have when looking at the cat could be had by me when perfectly hallucinating. So whether an experience is perceptual or hallucinatory has nothing to do with its intrinsic character. The intrinsic character of an experience is its phenomenal character; whether the experience is hallucinatory or perceptual depends on its causation.

In an influential series of articles Paul Snowdon, drawing upon earlier work by J. M. Hinton, presents a view of phenomenally visual experiences on which a difference in experience need not show up as a difference in the way things look.² Even if it is conceded that the way things look to me in a situation in which I see some scene is the same as the way things look to me when hallucinating, it does not follow that the experiences in the two situations are the same.

¹ It is not easy to spell out the nature of the required causal dependence. The problem is noted by H. P. Grice in ‘The Causal Theory of Perception’, *The Aristotelian Society, Supplementary Volume 35*, 1961.

² My exposition is based principally on Snowdon’s ‘Perception, Vision and Causation’, *Proceedings of the Aristotelian Society*, 81, 1980-81, 175-92 and ‘The Objects of Perceptual Experience’ *The Aristotelian Society, Supplementary Volume 64*, 1990, 121-50. Snowdon has further commented on disjunctivism in ‘What is Realism?’, *Proceedings of the Aristotelian Society*, 102, 2002, 201-228, and in ‘The Formulation of Disjunctivism: A Response

Fundamental to the disjunctivists' way of thinking is a *relational conception of perceptual experiences*. If I see a lemon in a basket on the kitchen table then under this conception I have an experience that is essentially relational because it is essentially a visual encounter with the lemon. If I were to have a hallucinatory counterpart of that experience—a perfectly hallucinatory experience that does not differ from the perceptual experience with respect the way things look—this would be a different experience because it is not essentially relational. It is not in dispute that ordinary perceptual experiences and perfectly hallucinatory experiences can be *phenomenally* the same, so that there is no difference between them with respect to the way things look to their subjects. The difference between traditionalists and *relationalists*, as we might call them, is over the conditions under which experiences are the same or different.

Snowdon presents disjunctivism as a view about the truth-conditions of sentences or judgements of the form 'It looks to S as if p' true. He writes

That 'looks' sentences are true in hallucinations and in perceptions, and are not ambiguous, does not entail that they are made true by (or true in virtue of) exactly the same conditions. ('The Objects of Perceptual Experience', p. 129.)

Omitting refinements, such sentences, and the judgements they can be used to express,

are made true by two types of occurrence: in hallucinations they are made true by some feature of a (non-object-involving) inner experience, whereas in perceptions they are made true by some feature of a certain relation to an object, a non-inner experience, (which does not involve such an inner experience). (*op. cit.* p. 130.)

Arguably, the broader picture with which disjunctivism is linked is separable from any specific concern to provide truth-conditions for looks-judgements, if that involves supplying necessary and sufficient conditions for the truth of these judgements. One might be a relationalist about perceptual experience while not engaging in any such enterprise. What, I think, is at the heart of

to Fish', *Proceedings of the Aristotelian Society*, 105, 2005, 129-41'. For Hinton's ideas, see his *Experiences* (Oxford: Clarendon Press, 1973).

disjunctivism is a view about what it is for it to be the case that it appears to one as if such-and-such is there. It captures the following ideas:

- (a) A case in which it appears to the subject (in some modality) just as if an F is there can be a case in which she has an essentially relational experience implicating an F or it can be a case in which this is not so.
- (b) Cases relating to (a) in which the subject is not having an essentially relational experience implicating an F include cases in which the subject is hallucinating an F and cases in which the subject perceives something which, though it appears to be an F, is not.
- (c) The case in which the subject has an essentially relational experience does not admit of decomposition into (i) the having of a non-relational experience (what Snowdon calls an inner experience), and (ii) the satisfaction of further conditions.

3. Perceptual knowledge, relationalism and the traditional conception

Snowdon's interest is not primarily epistemological. He is concerned with how we should conceive of perceptual experiences and he takes it that a constraint on an adequate conception is that it should make sense of how experiences enable us to think about objects demonstratively. When we perceive objects we are in a position to have demonstrative thoughts about them. Watching a tennis match I can think *That return was good*, where the reference of the demonstrative element of my thought, *that return*, is the very return I have just observed. Here I not only pick out perceptually a particular return, I am in a position to think of it, demonstratively, as *that return*. It is my having observed the return in question that enables me to do so. These considerations present a challenge to traditionalists about perceptual experience. If the experience that I have when I see an F is one I could have when not seeing an F, then it is not clear how the experience in the case in which I see an F enables me to think about the F in question demonstratively, never mind acquire knowledge about it. On this account it is not intrinsic to the experience that it presents an object to the subject. So there is explanatory work needed to show how it can contribute to making the object available as an object of

demonstrative thought. The attraction of the idea that perceptual experience is essentially relational is that it circumvents this problem. Just because the experience is essentially relational, the object is, so to speak, embraced by the experience and thus available to be thought about by suitably equipped subjects.³ That, at any rate, is the idea. It is clearly of interest for epistemology since any plausible view of perceptual knowledge must be consistent with facts about how perception makes demonstrative thoughts available to us.

In the work of John McDowell we find a disjunctivist style of thinking recruited to shed light directly on epistemological matters. Although McDowell acknowledges a debt to Snowdon and Hinton, it is not immediately obvious how the disjunctivist strand in his thinking links up with their distinctive concerns about the nature of perceptual experience. I shall first describe McDowell's epistemological thinking in its own terms before circling back to pick up the question of how it relates to the theory of perceptual experience.

In 'Criteria, Defeasibility, and Knowledge' McDowell criticises a way of thinking about knowledge that he thinks plays into the hands of sceptics.⁴ The scepticism that is the main target of the essay is not 'external world' scepticism but 'other minds' scepticism. The picture that is the main target is introduced in the following passage:

Judgements about other minds are, as a class, epistemologically problematic. Judgements about 'behaviour' and 'bodily' characteristics are, as a class, not epistemologically problematic; or at any rate, if they are, it is because of a different philosophical problem, which can be taken for these purposes to be separately dealt with. The challenge is to explain how our unproblematic intake of 'behavioural' and 'bodily' information can adequately warrant our problematic judgements about other minds. (*op. cit.*, p. 382)

³ See Snowdon's 'The Objects of Perceptual Experience', section VI. That demonstrative thought is philosophically important is emphasised by Strawson in *Individuals* (London, Methuen, 1959). The philosophers under consideration here have been influenced by the more developed discussion of the matter in Gareth Evans *The Varieties of Reference* (Oxford, Clarendon Press, 1982). The view is explored in McDowell's 'Singular Thought and the Extent of Inner Space' and figures crucially in Bill Brewer's *Perception and Reason* (Oxford, Clarendon Press, 1999) and John Campbell's *Reference and Attention* (Oxford, Clarendon Press, 2002).

⁴ The essay originally appeared in *Proceedings of the British Academy* 68, 1981, 455-79. Page references here are to the version in McDowell's *Meaning, Knowledge and Reality* (Cambridge, MA.: Harvard University Press, 1998), 369-394.

The picture incorporates an *evidentialist* model of knowledge of other minds. On this model judgements about other minds can be justified, and constitute knowledge, in virtue of being based on defeasible evidence relating to behavioural and bodily information that does not itself implicate mentalistic concepts. It is this evidentialist model that McDowell thinks plays into the hands of sceptics. For this purpose sceptics may be represented as follows.

- (a) They accept that if there were knowledge of other minds it would have to be on the lines of the evidentialist model.
- (b) They find the evidentialist model inadequate. Failing to appreciate that an alternative is available, they infer that we lack knowledge of other minds.

In McDowell's view sceptics are right to find the evidentialist model inadequate. One reason for thinking it inadequate concerns our entitlement to take facts about bodily behaviour, conceived in non-mentalistic terms, as evidence for the obtaining of psychological states that are not perceptually manifest. Another concern is focussed on the very idea that an adequate theory could 'envisage ascribing knowledge on the strength of something compatible with the falsity of what is supposedly known' (*op. cit.*, p. 372). This latter theme is not pursued in depth in 'Criteria, Defeasibility and Knowledge' though it is always in the background. (I shall defer discussion of it until section 5.) The main strategy in that essay has the following three stages:

The diagnosis

A version of the argument from illusion is viewed as motivating the evidentialist model of knowledge of other minds.

The way out

It is argued that the argument from illusion is not compelling. This is done with the help of a disjunctivist account of appearances and of a *perceptualist* model of some of our knowledge of other minds.

The response to the sceptic

The availability of a plausible perceptualist model is taken to show that the sceptic's argument is not compelling.

What is on offer, then, is an attempt to undermine the sceptic's argument. I turn now to a closer look at the diagnostic stage and the way out.

The argument from illusion, under McDowell's conception, runs as follows:

... since there can be deceptive cases experientially indistinguishable from non-deceptive cases, one's experiential intake—what one embraces within the scope of one's consciousness—must be the same in both kinds of case. In a deceptive case, one's experiential intake—what one embraces within the scope of one's consciousness—must be the same in both kinds of case. In a deceptive case, one's experiential intake must *ex hypothesi* fall short of the fact itself, in the sense of being consistent with there being no such fact. So that must be true, according to the argument, in a non-deceptive case too. One's capacity is a capacity to tell by looking: that is, on the basis of experiential intake. And even when this capacity does yield knowledge, we have to conceive of the basis as a *highest common factor* of what is available to experience in deceptive and non-deceptive cases alike, and hence as something that is a defeasible ground for the knowledge, though available with a certainty independent of whatever might put the knowledge in doubt.

...

In a deceptive case, what is embraced within the scope of experience is an appearance that such-and-such is the case, falling short of the fact: a *mere* appearance. So what is experienced in a non-deceptive case is a mere appearance too. (*op. cit.*, p. 386.)⁵

⁵ In its classic form the argument from illusion is not directly epistemological and not about which facts experiences 'take in'. It is about the objects of perception, in the sense of 'object' in which a thing is the object of perception when it is perceived.

The application to other minds cases is straightforward. Take the non-deceptive case—the good case—to be one in which Bill is anxious and looks it, and the deceptive case—the bad case—to be one in which Bill looks just the same as he does in the other case but is not anxious. In the two cases, according to the argument, what the experience takes in must be the same—Bill’s appearing a certain way that could be specified entirely in non-mentalistic terms. But assuming that what experience takes in must supply the justification of a judgement that Bill is anxious, the justification must be the same in both cases. It comprises the fact that Bill appears as he does and so is at best defeasible evidence that he is anxious. It is defeasible because it is possible that its evidential value should be undermined by further information, for instance, information to the effect that Bill is fooling around or maliciously deceiving. This is all in keeping with the evidentialist model.

The crucial inference in the argument from illusion, as conceived by McDowell here, is that

... since there can be deceptive cases experientially indistinguishable from non-deceptive cases, one’s experiential intake—what one embraces within the scope of one’s consciousness—must be the same in both kinds of case.

The good and the bad cases under consideration are experientially indistinguishable in the sense that in both cases how Bill looks is the same. In one sense, then, the common factor in the two cases is Bill’s visually appearing some way—a worldly appearance as we could call it. The conclusion drawn is that what the subject takes in must be the same in these two cases. What is taken in is what it seen to be so. That, according to the argument, has to be the same in the two cases, which is why the only plausible model for knowledge of other minds is an evidentialist one. The basis for a judgement to the effect that Bill is anxious has to be something that is as much available in the bad case as in the good case.⁶

Does McDowell deny the legitimacy of the crucial inference or does he reject the assumption of the inference? On the face of it, the right answer is that he challenges the legitimacy of the inference. (See the middle paragraph of *op. cit.* p. 389.) He thinks that *if* by ‘experiential indistinguishability’ you just mean that there is no difference with respect to how things appear

⁶ See Williamson *Knowledge and Its Limits* (Oxford: Oxford University Press, 2000), ch. 8 for further pertinent reflections on evidence in relation to pairs of good and bad cases.

then the good case and the bad case are experientially indistinguishable. But we are supposed to resist the conclusion that what the subject takes in is the same in both cases. This takes us to the stage of McDowell's strategy that constitutes what I called *the way out*.

The case has two complementary strands. The first is a disjunctivist conception of what it is for it to appear to a subject that such-and-such:

... an appearance that such-and-such is the case can be *either* a mere appearance *or* the fact that such-and-such is the case making itself perceptually manifest to someone. (*op. cit.*, pp. 386-87)

Notice that here appearances are not worldly appearances but subjective appearances. The focus is not on how Bill looks, but on how it looks to a subject that Bill is. (I do not think it distorts McDowell's position to think of subjective appearances in terms of how it appears to a subject *as if* things are.) We are to think of its looking to me as if Bill is anxious in these terms:

(a) Either it is a case in which the fact that Bill is anxious makes itself visibly manifest to me or it is a case in which Bill merely looks to me to be anxious, though he is not anxious.

(b) The case in which the fact that Bill is anxious makes itself visibly manifest is not decomposable into my having visible evidence that would be as much available in a corresponding bad case and making a judgement on that basis.

It is noteworthy that thus far there is no reason to think of McDowell as committed to the kind of disjunctivism about perceptual experience considered by Snowdon. The rejection of a common factor in the good and bad cases is not presented as the rejection of the idea that there is an experience common to the two cases. It is presented as a rejection of the idea that what the experiences take in must be the same, and along with that a rejection of the idea there is a common evidential base in the two cases. The disjunctivism exemplified by (a) and (b) above allows McDowell to concede that it can be true that it appears to me as if Bill is anxious in both the good case and the bad case and, indeed, that how he appears is the same in the two cases, *even though what the experiences takes in is different in the two cases*. For in the good case,

what is taken in is that Bill *is* anxious, and in the bad case what is taken in is merely that he has a certain look. But it is not clear at this stage whether the difference with respect to what the experiences take in requires there to be a difference in kind between the experiences.

Taken on its own this disjunctivist manoeuvre could easily seem like a sleight of hand. For it provides no explanation for *how it is* that what is taken in by an experience gained by looking can be different in the two cases. If there is no sleight of hand here it is because the disjunctivist manoeuvre is not supposed to carry the whole weight of the case that is designed to undermine the evidentialist model. It is supplemented by a perceptualist model of some of our knowledge of other minds.

Initially it might not seem promising to treat any knowledge of other minds as perceptual. It might seem bizarre that anyone should suppose that we can see straight off that someone is anxious when it is agreed on all sides that a person can display the demeanour of one who is anxious yet not be anxious. What makes sense of the perceptualist model is a conception of *perceptual knowledge as non-inferential knowledge acquired by suitably equipped subjects from what they perceive*. Perceptual knowledge is the kind of knowledge we acquire when we tell from its look that a bird is a magpie or that a flower is an orchid. It is non-inferential knowledge in that it is not acquired by reasoning from prior assumptions. It is phenomenologically immediate in that what is known simply strikes the subject as being so. To be suitably equipped subjects need more than the repertoire of concepts required in order to entertain the content of the knowledge. They need to have an appropriate sensibility incorporating appropriate recognitional abilities. McDowell applies this general conception to cases that philosophical tradition has treated in evidentialist terms. For instance, he applies it to knowledge of what people are saying from hearing what they say. In this case a pre-requisite is obviously knowledge of the language of the speaker and it is entirely plausible in this case that no inference, in any ordinary sense, is involved. Yet there is no doubt that knowledge of what people say is species of knowledge of other minds.⁷

⁷ For discussion of hearing what people say see, McDowell's 'Anti-Realism and the Epistemology of Understanding' originally in Hermann Parret and Jacques Bouvresse (eds.) *Meaning and Understanding* (Berlin, Walter de Gruyter, 1981), 228-48. Reprinted in *Meaning, Knowledge, and Reality*, 314-343. For a discussion of related issues concerning the scope of perceptual knowledge, see my 'The Scope of Perceptual Knowledge' *Philosophy* 75, 2000, 73-88.

Bringing sensibilities into the picture helps to explain why somebody *with the appropriate abilities* can detect that Bill is anxious from his look when others lacking those abilities can't. But what we are most concerned to understand is how somebody who does have the requisite abilities could be in a better epistemic position in the good case than in the corresponding bad case. To address this problem we need a fuller account of recognitional abilities. In particular, we have to think of the abilities as individuated in part by the environment of the subject. At this point I sketch a way of thinking about the abilities that can be deployed to fill out the perceptualist model.⁸

I have the ability to tell by looking, in suitable conditions for observation, whether or not something is a robin. That ability depends on my being disposed to go through certain judgement-forming procedures in response to suitable prompts. Suppose, for instance, that somebody wants to know whether a bird in the garden is a robin. If having looked in the right direction I have appropriate visual experiences then, absent countervailing factors, I shall judge that the thing in question is a robin. (It's not that I consider what kind of experiences I have; the procedure takes me from having the experiences to a judgement.) If I have a different range of experiences, then absent countervailing factors, I shall judge that it is not a robin. If I don't have a good enough look then I shall suspend judgement either way. In some few cases I would suspend judgement if the thing in question were not close enough to visual type to tell either way. (It might be a particularly bedraggled robin). The judgement-forming procedures are such that they reliably yield true judgements. But having such procedures in my repertoire does not suffice for my having the ability to tell by looking, under suitable conditions for observation, whether or not something is a robin. The very same procedures could come into operation in an unusual environment in which the visually robin-like things one encounters are at least as likely to be skilfully made robots as robins. If I were in such an environment I would not have the ability to tell by looking at things in that environment whether or not they are robins. To count as having the ability I must be in a favourable environment; it must not be one in which things that look just like robins could easily not be. So strictly speaking my ability is indexed to my actual environment and environments suitably like it: it is an ability to tell by looking, under suitable conditions for observation, whether or not something *in such environments* is a robin.

⁸ A closely related account figures in my 'What the Disjunctivist is Right About' *Philosophy and Phenomenological Research* (1997).

When we think of people as having the ability to tell by looking whether or not something is an F, we ascribe to them an ability of the sort I have described. But here we come to an important twist. (i) An implementation of the judgement-forming procedure that is productive of judgements to the effect that a thing before one is an F can yield a false judgement even if the environment is one with respect to which the subject has an ability to tell whether or not an F is present. The procedure is one which in that environment reliably yields true judgements, but that is compatible with its yielding a false judgement if through bad luck what is visually picked out is not an F, but looks just like an F. (ii) The relevant recognitional ability will have been exercised only if the resulting judgement is true. The ability is an ability to tell, that is, come to know, by looking whether or not something is an F. So if the subject does not come to know one way or the other then the ability will not have been exercised. The exercise of the ability is the resultant of an implementation of the procedure that yields a judgement, the subject's being embedded in a suitable environment, and the subject's getting it right on the occasion in question.

I have been concerned with abilities that enable one under suitable conditions and in a suitable environment to tell *whether or not* something is so. Clearly we have recognitional abilities that enable us to tell by looking that something is so in some cases when it is, but which would not generally suffice for telling by looking that it is so when it is, or for telling by looking that it is not so when it is not. My ability to tell from his look that Bill is anxious is like that. We may suppose that Bill has a way of looking—a certain demeanour—such that when he displays that demeanour I can tell just by looking at him that he is anxious. Even so, I may not in general be able to tell just by looking *whether or not* he is anxious.⁹

We can now explain why it should be thought that on encountering anxious-looking Bill I can tell that he is anxious *despite the fact that it is possible that he should look just as he does and not be anxious*. For the conditions for my exercising an ability to tell by looking that he is anxious can be met: the environment in this case is my normal environment; the relevant judgement-forming procedure is reliable in that environment, since Bill would never, or only rarely, display the demeanour in question, in the kind of context in question, if he were not anxious; the procedure is implemented and yields a true judgement.

⁹ Even so, as noted earlier, the latter ability should be conceived in such a way that there can be undecidable cases, for instances cases in which what is seen is not true to visual type.

The conception of recognitional abilities in play here does not provide the materials for a reductive conceptual analysis of knowing-that. Perceptual-recognitional abilities are abilities to tell that something is so, through looking, hearing, and so on. To exercise such abilities is to come to know something by some mode of perception. I have suggested that these abilities rest on reliable judgement-forming procedures, but it is not part of the suggested view that we can build up to a conceptual analysis of the abilities from a prior conception of the procedures. It is our conceptions of the abilities that are in the driving seat; our grip on the procedures is in terms of those conceptions. For instance, the experiences that are of the right kind for judging that something is an F are just the one's that would lead someone with the appropriate ability to judge that it is an F. It is doubtful that the range of such experiences is specifiable independently in general terms available to those who have a mastery of the concept of the ability in question. The fact that the abilities are characterised as abilities to know does not undermine their explanatory power. It is not as if they are conceived as powers, but goodness knows what powers, to acquire knowledge. They are powers of which we have commonsense conceptions, and which reflection reveals to depend on reliable judgement-forming procedures and to be individuated in part in terms of environments.

The unavailability of conceptual analysis of recognitional abilities in non-epistemic terms is in keeping with the fact that our most basic engagement with perceptual modes of acquiring knowledge—our own and those of others—is at the level of factive notions like *telling that* (*finding out that*), by means of seeing that, hearing that, and so forth. We seem to have learned to apply these notions in a largely satisfactory way without any guidance from analytical accounts on which they would be built up from more basic notions. It is arguably in favour of McDowell's way of thinking that these factive notions have this role.¹⁰

With this picture in place we can make sense of the idea that I can take in that Bill is anxious by looking. At any rate, if taking this in is a matter of recognising that Bill is anxious by looking then I can do that. A crucial difference from standard approaches to epistemology is the rejection of the idea that I know that he is anxious on the basis of evidence that is common to the good and bad cases. My knowing turns on my having exercised an ability to tell that Bill is anxious from his look, but that he looks this way is not to be conceived as evidence from which I infer that he

¹⁰ The same applies to the thinking of Timothy Williamson, which I have not been able to discuss here. See his *Knowledge and Its Limits*.

is anxious. This is just as well since I would be at a loss to sum up how Bill looks in the form of an assumption from which I could conclude that he is anxious, with the help of some suitable generalisation connecting the look so described with his being anxious. I simply tell by looking. This is in keeping with McDowell's *way out*. The other-minds sceptic is to be rebuffed by introducing a plausible perceptualist model of some knowledge of other minds. The model helps to make sense of McDowell's disjunctivism about appearances and in particular of his view that even though how it appears as if things are in the good and bad cases is the same the experiential intake of the subject is different in the two cases.

4. *The character of experience and what experience takes in*

As I noted earlier, Snowdon, like Hinton, is primarily concerned with the character of perceptual experiences, not with issues about knowledge.¹¹ Accepting the kind of disjunctivism I have already identified in 'Criteria, Defeasibility, and Knowledge' does not in any obvious way commit us to disjunctivism about perceptual experience. In the light of this it might be wondered whether McDowell need take a stance on the intrinsic character of perceptual experience. Other minds cases might suggest otherwise. It is not obvious that the view of recognitional abilities I have presented commits us to supposing that the experience I gain when I tell that Bill is anxious is different, *qua* experience, from the experience I would gain if Bill were feigning anxiety. Indeed, one might well take this to be in favour of the view since it is far from obvious that the visual experiences in the two cases need be different.

Nonetheless, McDowell does take a stance on the intrinsic character of perceptual experience. Such a stance is crucial for his view of what is needed to make sense of, what he calls, our vantage-point on the external world. Our vantage-point is threatened, he suggests, 'if subjectivity is confined to a tract of reality whose layout would be exactly as it is however things

¹¹ The same may be said of Michael Martin, who has developed disjunctivist thinking in ways that are akin to yet differ from that of Snowdon's disjunctivists. See Martin's 'The Reality of Appearances' in Mark Sainsbury (ed.) *Thought and Ontology* (Milan, FancoAngeli, 1997), 81-106, and 'The Limits of Self-Awareness' *Philosophical Studies*, 2004, 37-89. Snowdon comments on the variety of disjunctivist thinking in 'The Formulation of Disjunctivism'.

stood outside it'.¹² The threat comes from a picture of subjectivity on which perception of the world is indirect, being mediated by denizens of inner space. McDowell not only rejects such a view, he wishes to replace it with a picture on which experience is essentially world-involving and thus essentially relational.¹³ It is at this point that it is natural to think of his epistemological concerns as linking up with Snowdon's. The thought is that no account perceptual knowledge will be adequate unless it explains how a suitably equipped subject who has a perceptual experience is put in a position to have demonstrative thoughts about some object. By the lights of relationalists about perceptual experience, traditionalists lack the resources to account for how experiences can make objects available for demonstrative thought. If the experience I have as I look at the lemon in the basket could have been had by me if I were perfectly hallucinating then having the experience does not suffice to explain what makes the lemon available to me as an object of demonstrative thought. Notice that this view invokes a conception of demonstrative *thoughts* as essentially relational, and has it that where such thoughts are made available by perception the implicated experiences must also be essentially relational. It is a further step still to suppose that experiences can be essentially takings-in of *facts* as McDowell often appears to think. Which facts one counts as having taken in is, plausibly, determined by which recognitional abilities one exercises. So there is space for a position according to which although *perceptual experiences* are essentially relational it is not essential to any perceptual experience that any recognitional ability has been exercised and any particular fact taken in.

Traditionalists agree that perception is a subject-world relation. They may agree that a theory of perceptual knowledge must account for the possibility of perceptual-demonstrative thoughts—demonstrative thoughts about objects perceptually picked out. What they will certainly dispute is that the only plausible explanation will assume that the experiences implicated in perceptions are essentially relational. That leads us to the question whether traditionalists have the resources to account for the availability through perception of demonstrative thoughts.

¹² 'Singular Thought and the Extent of Inner Space' in Philip Pettit and John McDowell (eds.) *Subject, Thought, and Context*. Oxford: Clarendon Press, 137-68 and reproduced in *Meaning, Knowledge, and Reality*, 228-59. The quotation is on p. 241 of the reprint.

¹³ This is the burden of the portions of 'Singular Thought and the Extent of Inner Space' that deal with experience. See especially sections 5-8. A similar theme is in play in McDowell's use of the metaphor of openness to the world. See *Mind and World*, especially Lectures 1, II and VI, and the *Afterward, Part I*, pp. 140-41.

One way to do so would be to view the sort of recognitional abilities of which I have been speaking as depending on *abilities for perceptual discrimination*. I have in mind here not the ability to discriminate one type of object from another at the level of judgement, but rather a lower-level ability to pick out and track objects perceptually.¹⁴ Think of walking along a corridor encountering various potential obstacles. You pick them out visually. What this amounts to is that your visual experiences and dispositions to behaviour are shaped in a certain way. If, for instance, the people heading straight towards you were to continue on their current course you would move so as to avoid them. How you would move would be determined by other visual cues. Arguably, visual discrimination involves *sub-doxastic states* that prime subjects for behaviour, including *sub-intentional activity*. We engage in such activity when catching balls, reaching out to pick up objects, playing shots at tennis, and much else besides. It is plainly not the case that every movement you make when walking along a corridor is intentional even though the walking is. When you catch a ball in a game you do so intentionally but the systems that achieve hand-eye co-ordination result in small adjustments to the trajectory of your arm that are not intentional (or unintentional). Whereas the carrying out of intentions is guided by beliefs, the kind of sub-intentional activity of which I am speaking seems to be guided by changes in sub-doxastic state. The activity is responsive to cues that are picked up by sight but which need not be registered at the level of belief. Through the operation of the visual system, changes in experience and changes in behavioural dispositions vary systematically with changes in the environment and changes in the orientation and position of the subject, so as to enable the subject to move relatively smoothly in that environment. My picking out Bill by sight is not merely a matter of having a certain series of experiences but of having my experiences and my behavioural dispositions shaped as a result of the visual impact of Bill's presence upon me. It is a relational state targeted on Bill because it is his presence, and possibly changing location, that does the shaping and induces dispositions that enable me to be responsive to the relative positions of him and me in my egocentric space.

I see no reason why traditionalists about experience should not draw upon a view of visual discrimination along these lines in order to make sense of how perception makes objects available for demonstrative thought. Traditionalists may concede that an adequate account of perceptual knowledge must address the issue of how perception can enable us have essentially

¹⁴ I draw on such a view in 'What the Disjunctivist is Right About'.

relational demonstrative thoughts about physical objects. But it is not clear why they need take this to commit us to an essentially relational view of perceptual experience. Perceptual experience does enable us to pick out objects perceptually. But what the picking out consists in does not seem to require that the experiences themselves be essentially relational. Relationality comes in through the shaping by the object of the subject's experience and behavioural dispositions.

5. *Knowledge, justification, and reasons*

A central preoccupation of McDowell's *Mind and World*¹⁵ is with understanding how our thinking can be *rationally* constrained by perceptual experience—how it can be empirically grounded. His view is that to understand this we need to understand how perceptual experience can bear on our justification for believing this and that. To do that he thinks we must settle on the right way to think about how experiences give us cognitive access to facts. This in itself hardly seems controversial, for no one is going to disagree that empirical knowledge depends on perceiving things to be thus and thus. As McDowell sees it though, the experiences themselves must be essentially relational. Not only are they essentially pickings-out of objects, they are also essentially takings-in of facts about these objects. This is open to challenge even granted that perception furnishes us with knowledge. We surely do see, and thereby come to know, that this or that is so, and likewise for other modes of perception. But so far as I can see, here too, it is an open question whether such a view demands a conception of experience as essentially relational.

We can approach this strand in McDowell's thinking by taking stock of his treatment of an outstanding issue from 'Criteria, Defeasibility and Knowledge' concerning why it is not possible to explain knowledge in terms of the possession of evidence that is compatible with the falsity of what is known. In that article McDowell looked askance at theories that represent knowledge of other minds in these terms. For a developed view on what is supposed to be wrong with such theories we need to turn to 'Knowledge and the Internal'.¹⁶ The target of this article is what McDowell calls the interiorisation of the space of reasons.

¹⁵ (Cambridge, MA, Harvard University Press, 1984).

¹⁶ *Philosophy and Phenomenological Research* 55, 1995, 877-93. Page references are to the reprint in *Meaning, Knowledge and Reality*.

Our thinking operates in the space of reasons in that it is responsive to what there is reason for us to think and do. The interiorisation of this space that McDowell envisages is the upshot of a variant of the argument from illusion. Consider a case—the good case—in which I know because I see that, say, Mary has arrived. My knowing involves my taking it that Mary has arrived. By the argument, my taking this to be so is based on an experience such that it looks to me as if she has arrived. But it is possible, even if unlikely as things are, that I should have just such an experience, and believe accordingly, even though Mary has not arrived. This would be the corresponding bad case. In the bad case, by the argument, I have the same basis for thinking that Mary has arrived as in the good case. My standing in the space of reasons is therefore the same in the two cases. That standing has to do with what can be imputed to me as a rational thinker. It cannot turn on whether Mary has or has not arrived. It depends on whether or not my belief is justified. Under the prevailing assumptions, I am justified, and indeed have the same justification, in both the good and bad cases. To take such a view is to interiorise the space of reasons.¹⁷ It straightforwardly follows that knowledge itself cannot be, as McDowell puts it, a standing in the space of reasons. Knowledge depends on the satisfaction of the truth requirement and that is external to the space of reasons. The upshot on this way of thinking is that knowledge is to be conceived as a hybrid involving both a flawless standing in the (interiorised) space of reasons and the satisfaction the truth-requirement.

On a variant hybrid approach knowledge involves the operation of processes that reliably yield true beliefs plus the satisfaction of the truth-requirement. A belief might result from a reliable process and yet be false. So the truth-requirement is independent of the reliability requirement. On this type of hybrid view the very idea of the space of reasons plays no role so it is even more obvious that the acquisition of knowledge cannot be viewed as an achievement of a rational subject that is secured through operations in that space. The same applies even on a view that retains a role for justified belief but which also incorporates a reliability condition. For the satisfaction of the latter condition will be external to the space of reasons: it will depend on factors that are independent of the operations within that space. (For the themes of this paragraph, see *op. cit.* pp. 400-402).

¹⁷ I take the line of thought just set out to be the gist of the exposition of the argument from illusion in *op. cit.* p. 396.

It is bedrock for McDowell that knowledge itself is a satisfactory standing in the space of reasons. What is at issue is how this can be so. An important consideration for him is that it cannot be the case that the difference between a good and bad case, for instance, between the cases involving Mary just considered, consists simply in the fact that in the good case the truth-requirement is satisfied and in the bad case it is not. He thinks that it makes a difference to our standing in the space of reasons whether or not we have embraced a fact. For instance, it makes a difference to my standing in the space of reasons if I have *seen* that Mary has arrived as opposed to having a non-veridical experience such that it looks to me as if she has arrived.

Suppose that McDowell is right to reject the view that the only difference between the good and bad cases is that in the latter the belief is false and in the former it is true. Traditionalists about knowledge who adopt an interiorised conception of the space of reasons might well agree, for they may concede that the view in question makes it hard to capture the idea that knowledge is a matter of having a cognitive purchase on, or contact with, a fact. (For this language in McDowell, see *op. cit.*, p. 402.) For central cases of knowledge, one might think, there must be some connection between the subject's taking it that something is so and its being so. For cases of perceptual knowledge that *p*, a natural move is to require that the fact that *p* should have figured in the explanation of its looking to the subject as if *p*, and thereby coming to take it that *p*. It is open to defenders of the interiorised view of the space of reasons to take this plausible point on board. But they might also observe that it is unsurprising that there should be parity with respect to justification in good and bad cases: if justification does not guarantee truth there will pairs of good and bad cases, such that the justification is the same in each.

The requirement of an explanatory link between the fact that *p* and my taking it that *p* not only seems to be in keeping with the intuition that perceptual knowledge involves cognitive contact, it is required to address Gettier cases for perceptual knowledge. These are precisely cases of perceptually knowing that *p*, in which there is no appropriate explanatory link between the fact that *p* and the subject's taking it that *p*, but which are supposedly on a par with good cases not only with respect to justification, but also with respect to the satisfaction of the truth-requirement.

McDowell is committed to rejecting these ways of thinking about cognitive contact. The crucial consideration for him is that, on the proposed account, the explanatory link is external to the space of reasons. It does not account for what it is to have *cognitive* purchase on a fact just

because it does not display the posited explanatory link as relevant to accounting for the subject's achievement of an appropriate standing in the space of reasons. To explain that, according to McDowell, we have to reject the hybrid view. But is it so clear that knowledge can be such a standing? The view of perceptual-recognitional abilities introduced earlier enabled us to make some sense of how visual-perceptual knowledge about an object could be acquired on an occasion, despite the possibility that the object should look the same in a corresponding bad case. But nothing said in outlining the view makes clear how the knowledge supposedly acquired can meet the requirements for being a standing *in the space of reasons*. To address this matter we need to look more closely into the links between knowledge, justification, and reasons.

There are two aspects to the standing in the space of reasons in terms of which McDowell conceives of perceptual knowledge. One is that perceptual knowing is a state in which what is known is available to the subject so that it can figure as a reason, or an ingredient of a reason, for believing *other* things. The other more basic aspect is that emphasised in the following passage from 'Knowledge By Hearsay'.¹⁸

If knowledge is a standing in the space of reasons, someone whose taking things to be thus and so is a case of knowledge must have a reason (a justification) for taking things to be that way. But this is allowed if [for instance] remembering that Clinton is President is itself the relevant standing in the space of reasons. Someone who remembers that things are a certain way, like someone who sees that things are a certain way, has an excellent reason for taking it that things are that way; the excellence comes out in the fact that from the premise that one remembers that things are thus and so, as from the premise that one sees that things are thus and so, it follows that things *are* thus in so. (*op. cit.*, pp. 427-28.)

In this passage knowing that *p*, through seeing that *p*, is conceived as being in a state in which one has a reason to believe *that p* supplied by the fact that one sees that *p*. I shall call this *the basic point*.

To someone schooled in traditional approaches to epistemology the basic point could easily look like an abandonment of serious philosophical enquiry. It is hardly a matter of dispute that *if*

I have cognitive access to the fact that I see that p, then I have a justification for believing that p. So pointing this out will not by itself address the worry of someone who seeks illumination on how I can have cognitive access to such a thing. We are in the course of an attempt to shed light on the character of perceptual-cognitive access to facts and, in particular, on how such access can be an appropriate standing in the space of reasons. Addressing this issue by appeal to the availability of a fact, access to which is equally problematic, does not initially look like a satisfying response to the problem. Further explanation is surely needed. Nonetheless, the sheer commonsense of McDowell's thinking at this point should not be overlooked. In our ordinary thinking about knowledge and the possession of reasons we regularly treat ourselves as having reasons to believe something because we have seen it to be so. Epistemological theorising can make this seem problematic.

By reasoning similar to that which McDowell takes to be a species of the argument from illusion, many theorists will no doubt suppose that the only access I have to the fact that I see that p is via something that falls short of that fact—the fact that I have an experience such that it appears to me as if p. Just because this is of a piece with the thinking of arguments from illusion McDowell is committed to rejecting it. How then should we think of the access we have to facts as to what we see to be so? A natural step is to introduce recognitional abilities again. As we saw, being able to recognise by sight that an F is before one under certain circumstances is a matter of (i) having a grasp of the concept of an F, and (ii) having an ability to tell by sight under those circumstances that an F is present. We acquire the latter ability by learning when to apply the concept of an F in response to what we see. The second-order ability to apply to oneself the concept of seeing that an F is present is not so mysterious once we have an idea of the first-order ability to deploy the concept of an F recognitionally. The very same experiences that enable us to pick out an F and recognise it to be an F, enable us to tell that we see that an F is there, provided that we have the appropriate second-order recognitional abilities. These abilities are just as much indexed to a favourable environment as are the corresponding first-order abilities. If the environment is not right we may be primed to go through judgement-forming procedures that are needed for a second-order ability and yet lack the ability. But the environment can be right. The

¹⁸ This is in B. K. Matilal and A. Chakrabarti (eds.) *Knowing from Words: Western and Indian Philosophical Analyses of Understanding and Testimony* (Kluwer: Dordrecht, 1993), 195-224 and reprinted in *Meaning, Knowledge and Reality*, 414-43.

environment in which we can tell by looking that an F is there is also an environment in which we are in a position to know that we see that F is there.

With the notion of a second-order recognitional ability to hand the following looks like a plausible position:

- (a) We can know that an F is there in virtue of seeing (or having seen) an F and recognising (or having recognised) it to be an F.
- (b) When we know that an F is there in this way we have reason to believe that an F is there, constituted by the fact that we see (or saw) that an F is there.
- (c) We *have* the reason in question simply because the fact that constitutes it is available to us, thanks to our having an appropriate second-order recognitional ability—an ability to tell whether or not we see (or saw) that an F is (was) there.

It is not a requirement of the view that the second-order recognitional capacity should actually have been exercised on every occasion on which one knows. Even when it is not exercised there is a sense in which the corresponding belief is justified. In the usual case it will be justified since the fact that provides the justification is available and its availability will serve to keep the belief in place.

Something like the proposed view is required to make sense of the idea that knowledge is itself a standing in the space of reasons.¹⁹ It is crucial to appreciate that, as I have presented it, the view does not build up knowledge from justified belief. Rather the exercise of a recognitional ability explains how we know that p on a given occasion and since on such an occasion we can readily tell how we know—for instance by seeing that p—our believing that p can be justified. For all that I have said so far by way of outlining the view, there could be knowledge without justified belief. I have said nothing to rule out the possibility that there should be creatures with

¹⁹ John Greco takes McDowell to task for arguing in effect that content externalism is sufficient for rejection of scepticism. See ‘Externalism and Scepticism’ in R. Shantz (ed.) *The Externalist Challenge: New Studies on Cognition and Intentionality* (New York: De Gruyter, 2004). If the view I have been outlining is correct content externalism is not the only weapon in McDowell’s armoury against scepticism. There is, among other things, a model of justified belief at odds with the models that underpin scepticism.

knowledge who lack the second-order recognitional abilities in virtue of which facts as to what we see to be so are available to us. If knowledge is essentially a standing in the space of reasons then it can be so only for creatures with such abilities. So if creatures without such abilities can have knowledge, then knowledge is only contingently a standing in the space of reasons.

6. *A satisfying response to scepticism?*

McDowell from time to time appears to be dismissive of scepticism. He gives the impression that once we reinstate a commonsense picture on which we are open to the world, directly embracing worldly facts, scepticism ceases to be of interest.²⁰ This quietism underestimates the amount of headway that McDowell makes against certain sceptical arguments and, in particular, those that rest on the style of thinking that he calls the argument from illusion.

Sceptics have models of knowledge—conceptions of what would be required for knowledge of various kinds—and claim that the requirements for knowledge under their models are not satisfied. The challenge to sceptics that is discernible in the work of McDowell can be seen as having two components.

The first component is to challenge sceptics to defend their models of knowledge. Suppose that sceptics devise their models of knowledge, like everybody else, by considering cases that from a pre-theoretical standpoint would be judged to be central cases of knowledge and reflecting on why they are counted as cases of knowledge. Then the ground of battle is over whether sceptics are right to think that their models best account for the cases and for our grounds for thinking of them as cases of knowledge. From McDowell's perspective, as I understand it, sceptics lose this battle so far as perceptual knowledge and knowledge of other minds are concerned because their models of knowledge do not best account for the cases. The idea here is not that sceptics cannot do justice to the fact that the cases *are cases of knowledge*. That would obviously beg the question. It is that their models do not do justice to how in practice we think of knowledge. Suppose, however, that sceptics profess to adopt some different methodology for arriving at models of knowledge. Then we need to know what that methodology is. Otherwise the sceptic stands charged with giving us models of knowledge from nowhere. The trouble then would be that it's hard to see what sceptics can say. Scepticism will

have bite only if sceptics have correctly noticed requirements implicit in our actual thinking about knowledge that are never in practice satisfied. The anti-sceptical move at this stage—one that McDowell in effect makes—is to undermine the sceptical claim to have discerned such unsatisfiable requirements. That is the first component.

The second component is simply to point out that the requirements of knowledge are in fact often satisfied. To take this line is to assume something at issue in the dispute; it is to assume that we know this or that. But if we have the right models of knowledge we shall still have the dialectical advantage. For if sceptics have lost the battle over models for knowledge they are in no position to challenge the knowledge claims that in practice we take to be established in the ordinary run of things.

That in broad outline is the kind of challenge that can be mounted from the resources discernible in McDowell's thinking. It does not set out to give any special philosophical vindication of the claim that there are things we know. It takes on board the task of seeking an understanding of how we know what we profess to know, while showing that familiar truths about what we know, at least so far as we have been able to make sense of the knowledge in question, remain untouched by the moves made by sceptics.

The perspective I have suggested on the dialectic between sceptics and anti-sceptics might be questioned. In particular, it might be doubted that sceptics need make use of disputable models of knowledge rather than drawing attention to requirements implicit in our ordinary thinking about knowledge. There is a standard type of sceptical argument that might seem to do just that. The type is given by the following schema, where not-H is the negation of some sceptical hypothesis.

(a) If know that p then I am in a position to know that not-H.

(b) I am not in a position to know that not-H.

So,

(c) I do not know that p.

²⁰ See, for instance, *Mind and World*, p. 113.

Consider, for instance, a famous example discussed by Fred Dretske.²¹ At the zoo you see and recognise zebras in an enclosure. Consider, however, the sceptical hypothesis that the creatures in the enclosure are cleverly disguised mules. Do you have any reason to think this false? On the assumption that they would look just the same if they were cleverly disguised mules it might be thought that you have no reason to believe that they are not. But if that is right then it seems that you are not in a position to know that they are not disguised mules and in that case you do not know that they are zebras. As is well known some respond to this problem by denying that if you know that the creatures in the enclosure are zebras then you are in a position to know that they are not disguised mules. I share the view that this is not a plausible option.

On the view that I outlined earlier you may know that the creatures are zebras because you have a suitable recognitional ability indexed to the relevant environment and environments like it. This is not an environment that regularly throws up mules that are made to look like zebras. Moreover, you can have reason to think that the creatures are zebras because the fact that you have seen that they are zebras can be available to you, thanks to a second-order recognitional ability—an ability that is also indexed to your environment and those like it. Against this background sceptical arguments of the type under consideration can be seen to rely on a contentious assumption that does not drop out of commonsense thinking—the assumption that we do not know that the sceptical hypothesis is false. In the particular case in hand, you can know that it is false that the animals are disguised mules because you can see that they are zebras.

Crispin Wright presents a challenge to this approach.²² Suppose, for the sake of argument, that you know that the creatures are zebras and that you are justified (Wright says, ‘warranted’) in taking them to be. Does your justification transmit to the claim that they are not disguised mules? Surely not, Wright claims. Your justification, he thinks, is provided by the look of the beasts and has no bearing on the possibility that they are disguised mules. Moreover, if you did have justification for thinking that they are not disguised mules this would need to be independent of your recognition that their being zebras entails that they are not disguised mules.

²¹ ‘Epistemic Operators’, *Journal of Philosophy* 67, 1970, pp. 1007-23.

²² ‘(Anti-)Sceptics Simple and Subtle: G. E. Moore and John McDowell’, *Philosophy and Phenomenological Research*, 65, 2002, 330-48.

That is because the justification you have for thinking that they are zebras does not transmit to the claim that they are not disguised mules.

Wright's challenge undoubtedly has some intuitive appeal. At least when reflecting philosophically, it is not hard to put oneself in a frame of mind, in which it seems puzzling that you can be justified in thinking the beasts are not disguised mules when (a) all you have done is look at them and (b) it is possible that they should look as they do and be disguised mules. Nonetheless, Wright's challenge, as I have so far presented it, begs the question against the kind of view I have drawn from McDowell. It assumes that the justification for thinking that the animals are zebras does not discriminate between the state of affairs in which they are zebras and one in which they are disguised mules. That is to say it assumes that the justification would be the same even if the beasts had not been zebras. But on the view I have been taking seriously, the justification for thinking that the animals are zebras derives from your having *seen* that they are zebras. That does discriminate between the state of affairs in which they are zebras and that in which they are disguised mules. For if they had been disguised mules you would not have seen that they are zebras and would therefore not have had the justification you do for thinking that they are zebras. By contrast, if the justification for thinking the animals are zebras derived from an experience under the traditional conception, or from the fact that the creatures look just like zebras, it would not discriminate in the way required. (The experience would be same, and it would still be a fact that the animals looked just like zebras, in a state of affairs in which they were disguised mules.) The trouble now is that, while Wright's challenge begs the question, the view challenged seems only to have pushed the problem back onto the question what entitles you to think that you see that the animals are zebras. For (it might seem) the justification that you have for your thinking that you see that the animals are zebras must derive from how it appears to you as if things are, and that does not discriminate between the state of affairs in which you see that the animals are zebras and that in which you see animals that look just like zebras but are not. The point here, it should be noted, is independent of whether or not a disjunctive conception of appearances is correct. Even if the experience in virtue of which it appears to you just as if the animals are zebras is conceived as essentially one of embracing the fact that they are zebras, your justification for thinking that you are having *that experience* does not discriminate between the case in which you are having that experience and one in which you see animals that look just like zebras but aren't. So the argument goes.

Understandable though it is, this last move still begs the question. Wright's challenge assumes the following: that the good case in which you see zebras and the bad case in which you see disguised mules are on a par both with respect to justification of the belief that the animals are zebras, *and* with respect to justification of the belief that you *see* that they are zebras. But on McDowell's view there is no parity in either case. You have a reason to believe that the animals are zebras in the good case that you do not have in the bad case. The reason in the good case is that you see that the animals are zebras. Given the embellishment I suggested, we explain your having that reason in terms of the idea that you are in a position to tell that you see that the animals are zebras in virtue of your having an appropriate second-order recognitional ability. This last point, incidentally, explains why the view is not externalist about justification in any extreme sense. Let extreme externalism about justification be the view that you can be justified in believing something while not appreciating why you are justified. Then this view about perceptual knowledge and justification clearly does not fit the bill since the justification for believing that *p* when one sees that *p* is provided by the fact that one sees that *p* and that fact is one to which the subject has ready access.²³

The central issue, which Wright's discussion undoubtedly serves to bring out, is whether you have a reason to believe that the creatures are zebras in the good case that you do not have in the bad case. On the picture I envisage it suffices for you to know that the creatures are zebras that you have exercised a first-order ability to tell by looking that they are. What explains your acquiring knowledge on a given occasion is your having exercised such an ability. The ability is indexed to the sort of environment you are actually in and environments like it—it is an ability, when in such an environment to tell by looking at something in that environment that it is a zebra, when it is. In an unfavourable environment in which there is a decent chance that what look like zebras aren't, you do not have the ability, indexed to that unfavourable environment. In a favourable environment, but on an occasion in which, unusually, the animals before you are cleverly disguised mules, you have the ability indexed to that environment but do not exercise it when judging that the animals are zebras. To exercise it you would need to come to know that they are zebras, but obviously you didn't. On this picture knowledge is not built up from justified true belief. But rational agents such as ourselves, who have access to facts about what we think

²³ The relation between McDowell's views and internalist/externalist debates is explored in Ram Neta and Duncan Pritchard 'McDowell and the New Evil Genius' *Philosophy and Phenomenological Research*, forthcoming.

and what we see, have a reason for believing that p when we know that p through seeing that p . The reason consists in the fact that we see that p . This by itself will seem to skirt around the real issues if one insists that such facts are not readily available to us. But, arguably, they are readily available to us and through the exercise of second-order recognitional abilities. On this view knowledge-acquisition is possible only because various conditions are satisfied which we have done nothing to check out. That you know that the animals are zebras, and know that you see that they are, depends on your not being in a weird environment in which there is a decent chance that things looking like them are disguised mules. Given that you have the appropriate recognitional ability you need have done nothing independently of your having observed on the occasion in question that the animals are zebras, to satisfy yourself that you are not in such an environment. For some this latter claim will be the stumbling block. But it can reasonably seem to be a stumbling block only if it is assumed that knowing that the creatures are zebras would require prior assurance that one is not in a weird environment. That assumption is not obviously true. Indeed, if the model of perceptual knowledge in terms of recognitional abilities is correct then it is false.

Does this general approach help at all with radical scepticism about the external world, induced by instances of the sceptical argument schema above in which the sceptical hypothesis puts in question all knowledge of the external world? Arguably it does. Radical scepticism, as much as scepticism linked to less radical sceptical hypotheses, rests on models of knowledge that are disputable, since they impose requirements on knowing that it is not clear need be met. Because the models of knowledge are disputable they do not subvert the assumption that the cases we count as central cases of knowledge are actual cases of knowledge.