

4

The Content View

Recall the Inconsistent Triad of claims about the nature of perceptual experience and its objects that I set out in ch. 1.

- (I) Physical objects are mind-independent.
- (II) Physical objects are the direct objects of perception.
- (III) The direct objects of perception are mind-dependent.

Chs. 2 & 3 highlighted difficulties in rejecting (I) and (II) respectively. This leaves what is without doubt the most common strategy today: to reject (III). Chs. 4 and 5 consider two quite different variants of this strategy. The first, which is the subject of the present chapter, is close to current orthodoxy in one form or another; but I argue that it is unsatisfactory. The second, to which I turn in the following chapter, is far less widely acknowledged. Yet it is in my view superior precisely in remaining significantly closer to some of the core insights of the early modern period that in so many ways initiate the debate as I present it in this book. These have been lost in the perfectly reasonable contemporary rush away from Berkeley's and Locke's own responses of rejecting (I) and (II) respectively. The primary purpose of my whole discussion is to elaborate, motivate and defend this latter view. First, though, I consider the more familiar former variant of the strategy.

The leading idea is that perceptual experience is most fundamentally to be characterized by its representational content, roughly, by the way it represents things as being in the world around the perceiver. I call this the content view (CV). This approach has no place for the technical early modern notion of a direct object of experience as I have been using it. According to that early modern idea, perception is most fundamentally to be construed as the subject's acquaintance with specific such direct objects. This is what provides the most basic characterization of which experiential condition in question. According to (CV), on the other hand, the most fundamental account of our perceptual relation with the physical world is to be given in terms of the complete representational contents of perceptual experience rather than in terms of our relation with any kind of object. Thus (CV) rejects (III). For there are no direct objects in the early modern sense.¹ Likewise (II) is strictly speaking false according to (CV), although it is absolutely standard to reinstate another sense in which physical objects are nevertheless genuine objects of perception. They are the very things that we see, touch and so on, and not by seeing or touching anything else first either.² Correlatively, as we shall see in ch. 5 below, my own preferred

¹ I assume that it is a necessary condition on the truth of 'The F's are G' that there be F's.

² So proponents of (CV) may say that mind-independent physical objects are the 'direct objects' of perception according to their view. What they mean by this is that the account that they offer of the way in which we are presented with mind-independent physical objects has no need of any direct relation of acquaintance with mind-dependent entities of any kind. This is fine so far as it goes, although later in the present chapter I argue that their attempts to provide an adequate account of perceptual presentation in this context are unsuccessful. My point here is more straightforward. Throughout this book I use the expression 'direct object' to refer to those entities, if any, our direct acquaintance with which constitutes the most fundamental account of the nature of our perceptual experience. According to (CV) there are no such things. For the defining feature of their position is that the nature of perceptual experience is to be given directly in terms of its representational content instead.

The Content View

alternative strategy for rejecting (III) certainly must and can provide an account of facts about a person's perceptual relation with the physical world that are expressed in terms of something very like representational content, that a particular object looks thus and so to her, for example. The essential difference is that this object view (OV) insists, whereas (CV) denies, that such facts can only properly be understood on the basis of a more basic acquaintance relation between the subject and the objects of perception themselves very much in keeping with the early modern notion of a direct object.³

The obvious model of representational content for expounding (CV) is that of a person's thought about the world around him, as this is expressed in his linguistic communication with others, and registered by their everyday attitude ascriptions to him. Let us begin, then, with S's thought that a is E: a thought about a particular object in his environment, a, to the effect that it is E. Call this the initial model for content, (IM).⁴

The basic objection that I will be mounting over the course of this chapter is that, in assimilating perception to thought in this way, (CV) fails adequately to account for the fundamental difference between perception and thought: perception is an

³ Representative proponents of (CV) include Anscombe (1962), Armstrong (1968, ch. 10), Dretske (1981, ch. 6), Searle (1983, ch. 2), Burge (1986), Peacocke (1989, 1992, ch. 3), Harman (1990), McDowell (1986, 1994), Tye (1992, 1995, 2000), Byrne (2001). Although this is, as I say, close to current orthodoxy, there are other strong voices in opposition. These include Campbell (2002a, 2002b), Martin (2002), Travis (2004), Gupta (2006a, 2006b) and Johnston (2006).

⁴ My suggestion is not that (CV) simply identifies perception with thought. Indeed I set out below a number of the dimensions along which its various proponents explain the differences between these two key mental phenomena. It is rather that (IM) provides the canonical context for the introduction of the key technical notion of representational content that plays the fundamental role in the (CV) account of the nature of perceptual experience.

The Content View

experiential presentation of the physical world around us that makes an essential contribution to our most basic conception of what mind-independent physical objects are that is in turn crucial to our capacity for any genuine thought about particular such things and to our growing empirical knowledge about them.

Recall to begin with that (CV) holds that the most fundamental account of our perceptual relation with the physical world is to be given in terms of the complete representational contents of perceptual experience rather than in terms of our acquaintance relation with any kind of object: (CV) gives no role to the early modern notion of a direct object that provides the most basic characterization of which experiential condition in question. So it seems that there are just two directions for the development of (IM) in spelling out any particular variant of the position in such a way as to do justice to the difference between thought and perception.

First, it may be said that, although perception involves the very same contents as thought, the attitude is in this case different: perception is a matter of perceiving, as opposed to thinking, that a is F. At the present point in the dialectic, and in the absence of any further explication of the difference in attitudes involved, this is quite a weak move. For it consists in little more than a formal acknowledgement that perception is in certain respects like, and in other respects unlike, thought, without any substantive further theoretical illumination of the relevant similarities and differences. One possible strategy for developing the suggestion would be to appeal to the distinctive functional role of the perceptual propositional attitude, on the one hand, as against those of the other cognitive propositional attitudes. Still, the proposal faces an objection of principle when it comes to explaining the role of perception in

grounding the possibility of empirical thought. For perceiving that a is F either presupposes grasp of the content that a is F or it does not. If it does, then it is difficult to see what explanatory role perceiving that a is F might play in relation to the possibility of thinking that a is F, at least on the plausible assumption that grasping the content is at least a matter of already being able to think that thought. If it does not, then it is equally unclear how perceiving that a is F might be supposed to resolve any difficulty we may have in understanding how it is possible to think that a is F, since the latter clearly does require grasp of that very content. In the first case, any explanation presupposes what it is supposed to explain. In the second case, no explanation is obviously adequate; for grasp of the content that a is F is unnecessary for perceiving that a is F, so perceiving that a is F is clearly insufficient for grasp of that content.⁵ It would be possible at this point to deny that perception has any special role in grounding our capacity for thought about the mind-independent physical world around us, and hence to reject the need for any kind of explanation of this purported role. This is obviously a substantive debate that I cannot resolve here other than by explicitly acknowledging my own commitment to the broadly empiricist idea that perception does play a fundamental role in making empirical thought possible. In any case, without further theoretical structure in the distinction between the attitudes of perceiving and thinking, and in particular without an account of the interaction between the natures of these attitudes and the contents that are ‘available’ to them, this move is at best a preliminary to further substantive work.

⁵ This objection is a central focus of John Campbell’s case against (CV) (2002a, 2002b, ch. 6); and I read him as deploying it more generally also against variants of the view along my second direction for the development of (IM) that follows. A familiar reply insists that perceiving that a is F is matter of being related to a non-conceptual content that plays a crucial necessary but insufficient role in the complete account of grasping the conceptual content that a is F. See n. 7 below for the notion of non-conceptual content. Peacocke (1992) offers the most developed version of this strategy.

Positions along the second direction for the development of (IM) in giving a substantive (CV) account of our perceptual relation with the physical world focus directly on the contents of perception themselves.⁶ The basic idea is that perceptual contents are distinct from the paradigm contents of thought in ways that illuminate and explain the crucial theoretical differences between the two modes of mind involved. Since this is the primary focus of the vast majority of contemporary work in the area, I frame my own discussion below in these terms. Having said that, I believe that my objections to (CV) apply equally to variants along the first direction of development outlined above.

There are a number of dimensions along which differences between perceptual and thought contents may be marked. First, it is often said that, whereas the content of thought is essentially conceptual, the content of perception may be non-conceptual in nature.⁷ Second, it is also remarked that the contents of perception are most aptly characterized or expressed in demonstrative terms, without invoking specific, more arbitrary, conventional linguistic terms and categories; and perhaps that such demonstrative characterization applies both in connection with the singular and the predicational components of the content in question. So, whereas one might think, for example, that John is 6' tall or that Elly's dress is teal, one might instead and more directly express the content of one's experience when faced with the relevant objects

⁶ The two strategies may of course also be combined; but I cannot see how this possibility offers any further defence for (CV) against the objections given below.

⁷ Although the precise characterization of this distinction is highly controversial, it has generated a huge literature. For pioneering discussions sympathetic to the idea, see Evans (1982, esp. ch. 6), Peacocke (1992, ch. 3), Cussins (1990) and Crane (1992). Critical discussion may be found in McDowell (1994) and Brewer (1999, ch. 5; 2005). See Heck (2000), Kelly (2001), Peacocke (2001) and Byrne (2005) for non-conceptualist rejoinder.

in perception, as that that (man) is that tall (indicating or attending to his height) or that that (dress) is coloured thus (similarly indicating or attending to its colour).⁸

Third, philosophers sometimes note the passivity of perception in comparison with thought. Whereas one has a certain choice or freedom with respect to which contents occur in thought, one is in perception simply ‘saddled’ with determinate content.⁹

Fourth, it is also pointed out that perception has a degree of belief-independence not generally present in thought. If I judge that *p*, or indeed, on one fairly standard use of the term, think that *p*, then I believe that *p* and am in that sense committed with respect to *p*. There are cases of perception, in contrast, that suggest that a parallel commitment at least need not hold. When I see the Müller-Lyer illusion, in full knowledge of its status as an illusion, the lines still look different in length even though I do not believe that they are different in length. Thus many conclude that the content of perception is what it is in general quite independently of the subject’s beliefs in connection with that content. We may very often go on to believe what we see, but we need not do so.¹⁰

These various ideas about the distinctive nature of perceptual content may of course also be combined and/or opposed in many ways. I do not myself endorse any of them,

⁸ See McDowell (1994, Lect. III, and Afterword, Part II, 1998, Lect. III) and Brewer (1999, chs. 5 & 6).

⁹ Again this point needs handling with great care. One does have a certain amount of control over the content of perception through the direction and focus of one’s concerns and attention. On the other hand, although I may choose to think about the distribution of primes as opposed to the Israeli atrocities in Gaza, for example, when I do so I cannot but arrive at the conclusion that there are infinitely many primes when that particular question arises, provided that I attend to the relevant considerations accordingly. Still, many argue that there is a degree of spontaneity or freedom in thought not present in perception. See McDowell (1994, Lectures I & II), Peacocke, (2006, 2009) O’Brien (2007, 2009) and Soteriou (2009).

¹⁰ See Sellars (1997), Evans (1982, pp. 123 ff.), McDowell (1994, pp. 60-3), Brewer (1999, esp. § 5.3.3) and Smith (2001). For dissent, see Stroud (2001).

although they do all seem to me to be responses to genuine features of perception.¹¹ I mention them mainly to illustrate the range of theoretical options open to proponents of (CV) and then to set them aside with the comment that I intend my discussion in what follows to be conducted so far as possible at a level of abstraction sufficiently great not to depend upon any of these further commitments of various specific versions of the view.

The plan for the remainder of the chapter is as follows. First, at the start of § 4.1, I explain the way in which phenomena of illusion and hallucination provide a primary motivation for (CV) in the current context. Indeed, it is often assumed that only a version of (CV) of some kind is equipped adequately to accommodate these undeniable phenomena. Second, in the main body of § 4.1, I contend that (CV) accounts of illusion in particular are problematic on their own terms. Third, I argue that these difficulties are due to more fundamental defects in the whole (CV) approach: the inevitable involvement in perceptual contents of both the possibility of falsehood (§ 4.2) and the generality of predication (§ 4.3). I complete the picture by arguing in ch. 5 below that my own alternative (OV) offers a superior account of illusion and hallucination. So the purported advantages of (CV) are themselves entirely illusory.

There are of course other motivations for (CV). I comment very briefly here on three of these before turning to illusion and hallucination.

¹¹ See my discussion of the Müller-Lyer illusion and ch. 5 below for a return to some of the underlying themes here.

The Content View

First, it is sometimes said that only by thinking of perception in terms of its representational content is it possible adequately to account for the epistemological role of perceptual experience in connection with our empirical beliefs; and this point may also be combined with the claim that some such epistemological role is in turn essential to the very possibility of empirical belief about how things are in the mind-independent physical world around us.¹² I explain in ch. 6 below that (OV) is well placed simply to take over all that is correct in the positive (CV) account of the epistemological role of perceptual experience in connection with empirical belief; and indeed that it provides the indispensable foundation for that very account.

Second, arguments are given for (CV) from the phenomenon of perceptual transparency.¹³ The premise, very roughly, is that in attempting introspectively to scrutinize the nature of our perceptual experience we seem to alight directly upon the mind-independent physical world – at least as it appears to be – rather than any evident constituents or qualities of the experience itself. Hence it is supposed to follow that the nature of perceptual experience is to be given by how things appear in that experience to be in the mind-independent physical world, that is, by its representational content. The transparency premise is in my view quite compelling, although hardly straightforward to state precisely. Again, though, I claim that (OV) is at least as well placed as (CV) to endorse it. For the whole point of (OV) is to insist that mind-independent physical objects themselves are the direct objects of perceptual experience in my early modern sense. Indeed, it would be fair to say that most proponents of (CV) who appeal to the transparency claim in this way do so

¹² See e.g. McDowell (1994) and Brewer (1999).

¹³ See especially Tye (1995, 2000, 2002); and for helpful further discussion of the phenomenon of transparency itself, see Spener (in preparation (b)).

themselves precisely in order to motivate their position over some form of indirect realism theory along the lines discussed in ch. 3 above, or perhaps over versions of (CV) that appeal to non-representational qualia as well as worldly representational contents in accounting for the nature of perceptual experience.¹⁴ They may or may not assume that their position is the only possible alternative to such views. My point here is just that, regardless of the merits of the argument from transparency in that more limited context, it is not sufficient to motivate (CV) as against (OV).¹⁵

Third, arguments have been appearing recently that purport to establish something very like (CV) as a direct consequence of the very idea of (visual) perceptual experience of a mind-independent physical world.¹⁶ There are important differences between these arguments and their various proponents explicitly acknowledge that the dialectic is complex. Still, I focus briefly on Schellenberg's (forthcoming) 'Master Argument' in the hope of clarifying further the nature and content of my own opposition to (CV) as I intend it here. In fact I consider only the first half of Schellenberg's Master Argument as follows.

P1: If a subject is perceptually related to the world (and not suffering from blindsight etc.), then she is aware of the world.

P2: If a subject is aware of the world, then the world seems a certain way to her.

P3: If the world seems a certain way to her, then she has an experience with content *C*, where *C* corresponds to the way the world seems to her.

Conclusion 1: If a subject is perceptually related to the world (and not suffering from

¹⁴ For examples of the latter target see Peacocke (1983) and Block (1996, 1998 and 2003).

¹⁵ I return briefly to transparency at the end of ch. 6 below.

¹⁶ See especially Byrne (2009), Siegel (2010) and Schellenberg (forthcoming).

The Content View

blindsight etc.), then she has an experience with content *C*, where *C* corresponds to the way the world seems to her.

Even this segment of the argument raises many complex issues of both interpretation and philosophical substance. I offer an interpretation of the argument on which its conclusion is in genuine tension with my own opposition to (CV), although there are others on which this is not obviously the case; and I explain how and why I object to the argument on that interpretation. Thus, I take it that the idea of a person having an experience with content C, as this figures in both **P3** and **Conclusion 1**, is the idea of a person having an experience whose most fundamental nature is to be elucidated in terms of some kind of representational content *C*. **Conclusion 1** is therefore very close indeed to the defining thesis of (CV) as articulated in the opening pages of this chapter, that the most fundamental account of our perceptual relation with the physical world is to be given in terms of the representational contents of perceptual experience.

I assume an interpretation of what is involved in being aware of the world on which **P1** is totally unproblematic. In order to focus my opposition as precisely as possible, I also offer an interpretation of **P2** on which this is equally unobjectionable by my lights: if a subject is aware of the world, in vision, say, then there are truths of the form ‘o looks F’ that qualify in a perfectly natural sense to capture how things seem to the subject. In ch. 5 I distinguish two importantly different varieties of such claims; and there is a substantive question about whether any conjunction of such claims deserves to be regarded as the way things look. I leave both of these issues to one side

for present purposes, though, in order to clarify my principle objection, which may now be targeted specifically at **P3**.

I simply deny that it follows from the fact that there are truths of the form 'o looks F' that apply to a person S in virtue of her perceptual relation with o, that the most fundamental nature of that perceptual relation itself is to be characterized as having a perceptual experience with the representational content (of some kind) that o is F. On my interpretation of Schellenberg's Master Argument, this is the transition articulated explicitly by **P3**. The account of looks offered in ch. 5 below proves that **P3** is false on this interpretation. For I explain there precisely how various looks claims apply to S in virtue of her perceptual relation with the world around her without assuming that the very nature of that perceptual relation is itself to be characterized in terms of any corresponding worldly representational content. The perceptual relation between perceivers and the mind-independent physical objects in the world around them is on that account more basic than any such representational contents and grounds the truth of the looks claims that perfectly reasonably inspire talk of perceptual representation. For now, though, perhaps an analogy may be of some help in illustrating the tenor of my objection. If a mind independent object o exists, then there are certainly (perhaps indefinitely) many true sentences of the form 'o is F', but I would deny that o's existence itself consists in the truth of those sentences or can be fruitfully illuminated by listing the facts that o is F₁, o is F₂, ..., o is F_i, etc. O's existence is more basic than any such facts and is what grounds the truth of all those sentences. Similarly, if S sees a mind-independent physical object o, then there are certainly (perhaps indefinitely) many true sentences of the form 'o looks F', but I would similarly deny that S's seeing o itself consists in the truth of those sentences or can be fruitfully illuminated

by listing the facts that \underline{o} looks \underline{E}_1 , \underline{o} looks \underline{E}_2 , ..., \underline{o} looks \underline{E}_i , etc, or the fact that it visually seems to \underline{S} that \underline{o} is \underline{E}_1 , \underline{o} is \underline{E}_2 , ..., \underline{o} is \underline{E}_i , etc. \underline{S} 's seeing \underline{o} , her perceptual experiential relation with that particular mind-independent physical object is more basic than any such facts and is what grounds the truth of all those sentences as I attempt explain in ch. 5.¹⁷

All three of these arguments for (CV) raise major issues that I cannot address fully here, although I do insist that none of them conclusively favours (CV) over (OV).¹⁸ I focus primarily on illusion and hallucination in what follows for the following reason. Some version of (CV) may seem to be the only genuine option in the philosophy of perception if it is granted that the early modern conception of perception as conscious acquaintance with objects of some kind is bound to involve the postulation of mind-dependent such direct objects. For mind-dependent objects of experience or sense data of this kind are widely rejected today.¹⁹ As I outlined in ch. 1 above, the arguments from illusion and hallucination are highly influential in supporting precisely this conviction that the early modern approach to perception is indeed committed to mind-dependent direct objects of experience. My overall aim in this book, though, is to explain how the early modern insight that perception is most fundamentally to be construed as a matter of the subject's conscious acquaintance with certain direct objects of experience is essential to any adequate defence of

¹⁷ I claim that similar considerations apply equally against related arguments offered by Byrne (2009) and Siegel (2010).

¹⁸ At the same time I also acknowledge that some of the most sophisticated proponents regard the best case for (CV) as more like an inference to the best explanation of a whole range of disparate perceptual phenomena including those briefly discussed above and others besides. See especially Byrne (2001, 2009) and Pautz (forthcoming) for development of this idea.

¹⁹ Although I myself entirely endorse this rejection, it is certainly not universal. See, for example, Jackson (1977), Robinson (1994), Foster (2000) and O'Shaughnessy (2003).

empirical realism; and I argue in ch. 5 below that the identification of such direct objects with mind-independent physical objects themselves is perfectly compatible with the evident existence of illusion and hallucination. The point of the discussion to follow in §§ 4.1 and 4.2 is to demonstrate, somewhat ironically in this context, that (CV) has its own major difficulties in accounting for illusion that actually lead us back to precisely the early modern idea of perception as a relation of conscious acquaintance with mind-independent physical direct objects.

4.1 Illusion, Hallucination and Content

In an illusion a physical object, \underline{o} , looks \underline{E} , although \underline{o} is not actually \underline{E} .²⁰ Thus, according to (CV), an illusory perceptual experience is one with the false content that \underline{o} is \underline{E} . Similarly, in hallucination, it looks just as though there is a Φ -type physical object out there that is \underline{E} , when there is no such thing at all, or at least no such thing causally responsible in the right way for the experience in question; and (CV) simply characterizes this as a perceptual experience with the false content that ϕ is \underline{E} , for an appropriate empty name ‘ ϕ ’, or with the content that there is a Φ out there that is \underline{E} .²¹ Furthermore, given the arguments from illusion or hallucination as set out in ch. 1 above, it may seem as though the only alternative to this (CV) approach is to invoke mind-dependent direct objects of experience along the lines of Berkeley and Locke.

²⁰ As noted in ch. 1, there are visual illusions that do not meet this provisional characterization. See also Johnston (2006) and ch. 6 below. The condition may also fail to be sufficient. See ch. 1, n. 8.

²¹ In the case of veridical hallucination, this quantified content as it stands is true. The Φ out there is not seen, though, and the experience remains hallucinatory, since the Φ in question is not causally responsible in the right way for that very experience. Alternatively, following Searle (1983), it may be said that the relevant perceptual contents are causally self-referential: the Φ out there causally responsible (in the right way) for this experience is \underline{E} . This will be false in cases of ‘veridical hallucination’. Hallucination, like illusion, will in that case always involve false perceptual content.

This presents a choice between rejecting (I) and rejecting (II) from my opening Inconsistent Triad. I argue in chs. 2 and 3 above that neither is satisfactory; and this scepticism about the theoretical adequacy of appeals to mind-dependent direct objects of perception is in any case widely shared today. So it can seem very much as though (CV) offers the only possible approach to understanding our perceptual relation with the physical world.

In the present section I express some doubts directly concerning the (CV) account of illusion. In the following two sections I develop more extended objections to (CV) based upon two of the most fundamental characteristics of any appeal to perceptual content: the possibility of falsehood and the generality of predication.

Consider the Müller-Lyer illusion, (ML), for example, in which two lines that are actually identical in length are made to look different in length by the addition of misleading hashes. Rejecting any appeal to mind-dependent direct objects that actually do differ in length, as we surely must, the proponent of (CV) insists that we describe this as a case in which the lines are falsely represented in visual experience as being unequal in length: A is longer than B, say.²²

²² A possible (CV) alternative developed by Pautz (2009) is that both lines are truly represented as being of (equal) length, l , say; where the sense in which they nevertheless look different in length is that the perceiver has a strong inclination to believe that they are different in length. This seems to me quite an implausible approach. For the illusory look remains even for perceivers with no such inclination because they are fully aware of the illusory nature of the diagram. It is true that there may be perceivers with sufficiently unusual perceptual upbringing, as it were, for whom (ML) is not misleading: the lines do not look unequal in length. The current (CV) proposal that such lines are accurately represented as each of length l even by perceivers for whom they do misleadingly look unequal in length is absolutely not required to account for this possibility though. See § 5.3 below for my own (OV) alternative account.

The Content View

My concerns about this account revolve around the question how exactly the world must be for the (ML)-experience to be veridical. (CV)'s insistence on characterizing perception by its content requires a specific answer to this question. Yet it is far from clear how one is non-arbitrarily supposed to be given, or even what the parameters are for making progress towards such an answer.

To begin with, is the line with inward hashes supposed to be represented as shorter than it actually is; or is the line with outward hashes supposed to be represented as longer than it actually is; or both; and by how much in each case? The most minimal answer here is that the content is simply, and no more than, that A is longer than B.

This is McDowell's answer.

Suppose I say of two lines that are in fact the same length, that one, say A, is longer than the other, say B. In saying that, I am representing A as longer than B. It does not follow that I am saying that A is longer than it is, or saying that B is shorter than it is, or saying that both of these things are the case. One of those things would have to be so if what I say were true, but I am not saying of any one of them that it is so. A 'by how much?' question does not arise. Just so with an experience that represents one line as longer than the other. (2008a, p. 201)

It is fine so far as it goes; but this is surely not far enough. The idea, in the terms of (CV), that a normal perception of (ML) simply represents that A is longer than B, without offering any other information whatsoever about the distribution of the two lines in space, whether they are millimetres or miles long, whether they differ in length by very little or by a great deal, and so on, strikes me as wildly implausible.²³

For lines that we see before us just do look to be of certain specific lengths at least in standard conditions. Of course there are difficult questions for proponents of (CV)

²³ McDowell does admit that there will be far more to the content of any actual (ML)-experience than simply that A is longer than B; but the example he gives is of the colours that lines and background are represented as having (2008a, p. 201).

The Content View

and everyone else about how best to capture the way in which their lengths are given in perception; but simply to deny it seems to me disingenuous. This is precisely a salient difference between saying and seeing. I may say that A is longer than B without saying anything at all about how long either A or B is, or by how much, whereas I normally see that A is longer than B by seeing the extent of A and the extent of B and noticing the former is greater than the latter by roughly such and such an amount. At least that is how it seems to me; and some such intuition figures as an important part of the motivation for various appeals to non-conceptual content or essentially demonstrative content in (CV) explanations of the distinction between the contents of perception and thought. The idea in both cases is that perceptual experience in general displays a far greater degree of determinacy and fineness of grain than is normally present in thought. McDowell's proposal that, so far as their lengths are concerned, the content of our perceptual experience of (ML) lines represents these merely determinably as one longer than the other, without assigning either any determinate length or giving any determinacy to the degree of difference in their lengths is at odds with this whole way of thinking.

Matters are further complicated when we acknowledge that (CV) appears committed to some kind of conflict within the content of the (ML)-experience itself. For the endpoint of each line certainly looks to me to be exactly where it actually is, at its actual position on the (ML) diagram. (CV) presumably captures this fact with the claim that my perception has as part of its content that there are lines joining the relevant two pairs of endpoints: one joining a to a' (A) and one joining b to b' (B). At the same time, (CV) registers the illusory nature of my experience with the claim that it is also part of its representational content that A is longer than B. Such a conjoined

content is not straightforwardly contradictory, as McDowell again points out.²⁴ For it does not follow from the fact that “I say something that places four points determinately where they are in the objective world ... and the distance between one pair of points is the same as the distance between the other pair” that “I say that the distances are the same” (2008a, p. 201). Still, (CV) is thereby committed to regarding the representational content of my (ML)-experience as impossible: it cannot possibly be veridical. There is nothing obviously wrong with making or endorsing an impossible claim: ‘Hesperus is distinct from Phosphorus’ is plausibly such. It does strike me as unattractive, though, to have to admit that this is what perception is bound to do whenever faced with the (ML) diagram. This would not be analogous to viewing an ambiguous figure, such as the necker cube, with respect to which observers may switch between mutually inconsistent aspects. The situation according to (CV) would be rather that (ML) compels us perceptually to represent reality as being a way that it could not possibly be, in a single take, as it were.

Perhaps proponents of (CV) may simply accept this result with equanimity. It does seem to be in serious tension, though, with relatively widely accepted views about the relation between conceivability and possibility.²⁵ Nobody thinks that distinct perceptions, even perceptions of a single subject over a short period of time, need conjoin consistently or possibly. But the idea that what we can conceive is a guide to what is possible plausibly rests, at least in connection with the perceptible

²⁴ There are those who regard even contradictory perceptual contents as perfectly unproblematic. See Crane (1988, 1992), who gives this as an argument for the non-conceptual nature of perceptual content. If the considerations that I offer below against impossible perceptual contents are sounds, then this stronger commitment to the possibility of contradiction is also unacceptable. See § 5.3 for more on this issue.

²⁵ See Gendler and Hawthorne (2002) for significant recent contributions to the debate.

The Content View

macroscopic physical world, on a combination of the following two principles. First, that what we can conceive in connection with the macroscopic physical world is constrained by what we may experience in a single perception: very crudely, perceptual experience provides the raw materials of genuine conception in this domain and also controls the ways in which these may legitimately be employed. Second, that what we may experience in a single perception is at least a guide to what is possible: if we really are able to take something in in a single perception, then this must at least be a possible state of the world. Given these two principles, it follows that legitimate conception of this kind will be, as many believe that it is, a guide to what is possible with respect to the perceptible macroscopic physical world.

This is clearly a major topic, and there are no doubt numerous qualifications and refinements to add; but I confine myself at this stage to two salient points. First, this line of thought provides a prima facie obstacle to the (CV) proposal that (ML) and other similar illusions automatically and unavoidably produce impossible perceptual contents in a single experience. Second, the natural response in connection with (ML) to the idea that what we may experience in a single perception must at least be a possible state of the world is surely to distinguish in some way between our awareness of the lines themselves, distributed in space as they are between two pairs of equidistant points, on the one hand, from their misleading appearance as unequal in length, on the other. This is precisely the approach taken by my own preferred (OV) in ch. 5 below. (CV) is restricted to accommodate all that there is to the nature of a single perceptual experience in its representational content; and, as we have seen, this appears to undermine the link between conceivability and possibility by being an impossible content in the case of (ML) and other similar illusions.

A first response at this point on behalf of (CV) would be directly to question the claim that what we may experience in a single perception is at least a guide to what is possible. For we are familiar from the drawings of Escher with numerous perceptions of ‘impossible figures’.²⁶ We must be careful here, though, about the distinction between the way in which the lines of the drawing itself are perceptually represented as distributed on the page, on the one hand, and the way in which a diagram as so represented might in turn be taken to represent another state of affairs. Only the latter and not the former is impossible in Escher cases. Yet the perceptual representation that is held to be at least a guide to possibility is the former and not the latter. The difficulty for (CV) is that our experience of (ML) appears committed to impossibility in the former: in our most basic representation of the lines of the diagram itself. For these are represented both as extending between two equidistant pairs of points and as being unequal in length.

It may be replied instead, then, that (CV) is not absolutely committed to the idea that illusions like (ML) produce impossible perceptual contents in this way. One such move would be to claim that our perceptual content conjoins a representation of one line as longer than the other with a placement of their endpoints as lying somewhere in a region around their actual position. This conjunction has possible instantiations, in which A is a little longer than it actually is, but with endpoints still within the relevant regions and B is similarly a little shorter than it actually is, for example. This is not true to my own experience of good (ML) diagrams, though. For I at least seem to experience the locations of the endpoints of the lines as closer to their actual

²⁶ See Hofstadter (1979) for many examples and interesting discussion.

The Content View

positions than is compatible with the extent by which they simultaneously look to be different in length.

Another more promising move would be to insist upon the crucial role of attention in modulating our experience of (ML). Attending to the endpoints, these look to be just where they actually are, within a relatively small margin for error; attending instead to the relative lengths of the lines, these look to be significantly different. So there is no single experiential take in which the proponent of (CV) is committed to an impossible content. As I explain in ch. 5 below, there is something right in this idea that different, and possibly even incompatible, looks, or appearances, may be produced by shifts in attention to a single unchanging object. The difficulty for proponents of (CV) at this point in the argument, though, is to accommodate the fact that the constancy in length of the two (ML) lines is also absolutely evident in our experience of shifting attention from their endpoints to their relative lengths and back. Thus the lines that are represented as unequal in length are represented there and then as those that extend between the pairs of end points a-a' and b-b', and similarly vice versa: the lines that are represented as extending between endpoints a-a' and b-b' are represented there and then as those that are unequal in length. Since this whole combination is not possible, (CV) is indeed committed to impossible overall representational contents. The lesson I draw in my own (OV) is that there is a fundamental level of perceptual experience whose nature is not to be captured in terms of any representational content, but rather in terms of a conscious relation of acquaintance between the subject and the particular mind-independent objects in question. This grounds and explains all the various ways that those things may look. But this is precisely what (CV) cannot admit. For the defining feature of (CV) is commitment to the idea that

The Content View

the most fundamental account of our perceptual relation with the physical world is to be given in terms of the complete representational contents of perceptual experience rather than in terms of our relation with any kind of object. This is what leads to the problematic endorsement of impossible overall contents in certain cases of illusion.

The attempt on behalf of (CV) to invoke shifts in attention in order to avoid the postulation of impossible contents that are in tension with the role of perception-based imagination in the epistemology of possibility suggests to me at least both the fundamental source of the difficulty for (CV) and the direction in which to pursue a more satisfactory alternative. The proposal was that shifting attention between the locations of the endpoints of the main lines of the (ML) diagram, on the one hand, and their relative lengths, on the other, produces successive contents accurately representing the endpoint locations and illusorily representing the relative lengths. The conjunction may be impossible, but each on its own is perfectly possible. The latter representation of A as longer than B is simply false: hence the (ML) illusion. My response was that the constant extent of the two main lines is surely evident throughout our experience of the (ML) diagram, even whilst we are attending to their relative lengths and thus misrepresenting one as longer than the other. The only way that (CV) can register the experiential presence of that constant extent is by appeal to a content that represents those lines as extending between the accurately represented endpoints. Thus, the impossible conjunction shows up in a single experiential take. The obvious alternative that I develop at length in ch. 5 below is that the evident presence in experience throughout of the unchanging lines is to be captured, not in terms of representational content at all, but in terms of a conscious relation of acquaintance between the subject and the mind-independent object in question: the

(ML) diagram itself. In this way, I contend, the phenomenon of the presentation of mind-independent physical objects in perceptual experience is not a matter of a level of representational content that gets something about the object right, and which therefore opens the door to impossible contents in various cases of illusion in which the very same features of the object are according to (CV) also misrepresented. It is a direct experiential relation between the subject and the particular mind-independent object in question itself, that is the fundamental ground of the possibility of getting anything right or wrong in perception at all. The postulation of impossible contents in some illusions is therefore just a symptom of more fundamental problems at the heart of (CV). I turn now to a more direct discussion what I take these to be: the possibility of falsehood (§ 4.2) and the generality of predication (§ 4.3).

4.2 The Possibility of Falsehood

According to (CV), the most fundamental account of our perceptual relation with the physical world is to be given in terms of the representational content of perceptual experience. Such contents are determined as true or false by the way things actually are in the world around the perceiver. Various ways in which they may turn out to be false appear well suited to accommodate the evident possibilities of illusion and hallucination. In particular, in the case of illusion, a particular mind-independent physical object o looks F although it is not actually F. Thus, according to (CV), it is o that looks F because o is presented in perception; and this is presumably due to the fact that the content of the perception in question in some way concerns o. Yet o is not the way that that content represents it as being, namely F. This combination is perfectly unproblematic in thought; but when the errors become very significant it

The Content View

seems to me to put serious pressure upon the (CV) account of experiential presentation in perception. For this involves actually seeing or otherwise consciously perceiving that very object.

I begin with a toy example to illustrate the point. I can certainly think, of a figure that you hide behind a screen, that it is square, when actually it is circular; but, if we insist on characterizing my perceptual experience as a representation of something as square directly before me, then how can we claim that it is actually a circle that is thereby perceptually presented, even if there actually is a circle out there, where I represent a square as being, which is somehow causally relevant to my purported perceptual representation? Crudely, if all that (CV) has to go on in accounting for the phenomenon of perceptual presentation is the representational content of the experience in question, then this central notion of perceptual content seems to come under serious tension from demands that pull in opposite directions. On the one hand, the phenomenology of genuine perceptual presentation surely places certain limits on the nature and extent of any errors involved. On the other hand, the basic notion of false content, which is crucial to the (CV) account of illusion, appears subject to far less demanding, if any, such limits. I now attempt to spell out this line of argument against (CV) in more detail.

We have been agreed from the outset that physical objects are presented to us in perception. This is intended to capture the utterly uncontroversial sense in which we see and otherwise consciously perceive physical objects without commitment to any specific theoretical elucidation of what such perceptual presentation consists in.

According to (CV), the most fundamental account of our perceptual relation with the

physical world is to be given in terms of the representational content of perceptual experience. (CV) has no place for any more basic or independent notion of a direct object of acquaintance in perception. Hence the account of perceptual presentation must be given entirely in terms of content. Presumably the basic form of the account is that a mind-independent physical object \underline{o} is presented in perception just if the content of the experience in question concerns \underline{o} . I use the technical term ‘concerns’ here in order to remain neutral on the precise account of this relation to be given. For there are many different approaches taken to account for presentation along these lines.²⁷ One important question is whether the contents of perception are general or singular. In the former case a range of illustrative accounts would be as follows. A perceptual content of the form $\langle \Phi x, Fx \rangle$ concerns \underline{o} just if \underline{o} is (a) the unique Φ in some appropriately demarcated domain of the physical world around the perceiver, or (b) the unique closest approximator in some sense to be specified to a Φ in some appropriately demarcated domain of the physical world around the perceiver, or (c) the Φ in the relevant region of the physical world that is appropriately involved in the production of the experience in question, or (d) the closest approximator in some sense to be specified to a Φ in the relevant region of the physical world that is appropriately involved in the production of the experience in question, or...²⁸ In any case, the idea would be that the content is veridical just if the relevant \underline{o} is \underline{F} . In the latter case of singular perceptual contents the idea would be that a content of the form $\langle \underline{F}a \rangle$ concerns \underline{o} just if \underline{o} is the reference of the singular content component ‘ \underline{a} ’, and

²⁷ The literature here is immense. For some of the most important discussions of the way in which perceptual contents may concern particular mind-independent physical objects, in my sense, see Grice (1962), Lewis (1980), Dretske (1981), Searle (1983, 1991), Armstrong (1991), Burge (1991), McDowell (1991), Davies (1992), Owens (1992) and Soteriou (2000).

²⁸ There will surely be some causal requirement here. The difference between (a) and (b), on the one hand, and (c) and (d), on the other, is that the former build this into the content of Φ itself, whereas the latter do not.

The Content View

once again the content as a whole is veridical just if the relevant \underline{o} is \underline{F} . No doubt there are many other possibilities besides. The basic proposal, though, is that a physical object \underline{o} is presented in a perceptual experience just if its content concerns \underline{o} .

For present purposes, we may characterize a visual illusion as a perceptual experience in which a physical object, \underline{o} , looks \underline{F} , although \underline{o} is not actually \underline{F} .²⁹ It is an absolutely essential feature of illusion, as opposed to hallucination, that some object, \underline{o} , is indeed presented in the illusory experience. Thus, according to (CV), a visual illusion is a perceptual experience with a content that concerns \underline{o} , and represents it as \underline{F} although it is not in fact \underline{F} .

Now, it seems to me that there are limits on the nature and extent of any errors involved in illusion, as opposed to hallucination, due to the fact that this involves genuine perceptual presentation of the object \underline{o} that illusorily looks \underline{F} . These are not hard and fast, there may be wildly abnormal circumstances in which the illusions I reject are in fact possible, and I certainly don't propose strict principles governing which are and which are not possible, or even suppose that there are any such general principles. Still, although a pair of lines of equal length about a metre away directly in front of me may look unequal in length, or may look unparallel when they are parallel,³⁰ those very lines could not normally look like a single perfect circle, or two perfect circles, for example. Similarly, although a rabbit curled up on the chair next to me may look like a cat or a cushion, that very animal could not normally look like the Eiffel tower. This is due to the fact that genuine perceptual presentation is in general incompatible with extreme error. I suggest that this in turn is a consequence of the

²⁹ See above and ch. 6 for caveats to this provisional characterization.

³⁰ For a classic illusion involving the latter, see Zöllner (1860).

The Content View

fact that, just as perceptual presentation is the source of a rough and provisional conception at least of what physical objects are, it is likewise the source of a rough and provisional conception at least of what the particular presented physical objects in question are. This explanatory proposal is not crucial to the argument, though. The datum, I claim, is that there are limits on the nature and extent of any errors involved in illusion, as opposed to hallucination, due to the fact that the former, but not the latter, involves the presentation of particular physical objects in perceptual experience. The challenge to (CV) is to account for these limits, given that they are not obviously entailed by the structure of the position so far.

It may be objected right away that any such limits as there may be on the nature and extent of the perceptual errors that are compatible with illusion, in which a mind-independent physical object is genuinely presented in perception, as opposed to hallucination, in which no such object is presented, are merely contingent upon the nature of the environment and the workings of subjects' perceptual systems. They are not a consequence of the very nature of perceptual presentation itself and so cannot be exploited in arguing in general against (CV) as I propose. In reply I accept that there is a certain degree of contingency in the limits that I contend govern the kind of perceptual error that is compatible with genuine presentation. The range of ways that an object of a given nature may look, for example, certainly depends to an extent upon the peculiarities of the environment and perhaps also upon the match between this and the context in which subjects' perceptual systems evolved. Still, I claim, given a fixed environment and evolutionary context, there are limits beyond which an object fails to be genuinely presented in perception regardless of its causal involvement in the production of a representation with the relevant false content. To

The Content View

anticipate a little, I argue in ch. 5 below that the ways that an object looks are a function of its de facto visually relevant similarities, in the circumstances and from the point of view in question, with paradigms of various perceptible kinds. Which similarities are visually relevant may depend upon evolutionary and experiential history; and which such similarities an object actually has in given circumstances from a particular point of view clearly depends upon the relevant environment. Still, as I say, there are limits beyond which an object fails to be presented in perception regardless of its role in causing an experience in which there looks to be something thus and so out there. Thus, although variation in the environment or evolutionary context may modify these limits to some extent, some such limits remain in any setting. It is these that I contend (CV) struggles to accommodate. If the objector's claim is that anything may illusorily look any way in any context, then I simply disagree.

So the situation is this. (CV) accounts for the presentation of a mind-independent physical object o in experience in terms of the fact that the subject's perceptual content concerns o in one of the various ways outlined above. The (CV) account of illusion entails that such presentation must be compatible with error in the predicational component of the perceptual content in question. Given a certain environmental and evolutionary setting, there are limits on the nature and extent of such error compatible with genuine presentation, beyond which the subject's experience must be regarded as a hallucination of some kind in which o is not subjectively presented at all, regardless of the fact that it may be involved in the causal production of the perceptual representation in question. The challenge for (CV) is adequately to acknowledge, incorporate and account for such limits. For they are

not an obvious consequence of the basic notion of a representational content simply concerning a mind-independent physical object as elucidated so far. I consider general and singular perceptual contents in turn to illustrate the point and the problem for (CV).³¹

On the side of general perceptual contents, there is no restriction whatsoever, so far as I can see on the nature or extent of any error involved in a content of the form $\langle \Phi x, Fx \rangle$, unless this is accomplished by stealth, as it were, by Φ already imposing a condition that the object concerned, o , is not sufficiently different from an \underline{F} to prohibit illusory presentation as such. And I can see no reason to expect such stealth always and everywhere to obtain. To give a concrete example, there is no obvious bar, for all that has been said so far, to each of the following contents concerning an (ML) diagram: $\langle x$ is directly in front of me (and causing this experience), x consists of two lines, the upper longer than the lower \rangle ; $\langle x$ is directly in front of me (and causing this experience), x is a perfect circle \rangle . Yet the right result here is surely that, although the first may well be an illusory experience in which the (ML) lines look unequal in length, the second is at best a hallucination that may in some way be caused by the

³¹ An important (CV) response here would be to deny the need for an explanatory account of the kind that I am asking for of the limits on the nature and extent of the predicational errors that are compatible with genuine perceptual presentation. According to (CV), presentation depends upon a degree of matching between the representational content of perceptual experience and the nature of the objects that cause it; but it is no part of the position that any complete specification may be given of the details of this matching condition. Perhaps there are even reasons for thinking that it is in principle uncodifiable. This is certainly often said to be the case in connection with other structurally similar supervenience claims. Why might it not also be the case here? This protest clearly raises a number of major issues about the point and purpose of the theoretical project here and the basic assumptions on which it is based. The substance of my response is given in § 4.3 below in the course of a critical consideration of the proposal that being presented in perception with a particular mind-independent physical object might be a composite condition, with an internal experiential matching component, I, and an external environmental-cum-causal component, E.

The Content View

presence of the (ML) diagram. Only the former, and not the latter, involves the perceptual presentation of the very (ML) diagram itself before the subject. The generalist proponent of (CV) may mimic this result by revising the account of perceptual presentation along the following lines. A content of the form $\langle \Phi x, Fx \rangle$ presents o just if it concerns o , as above, and o is not beyond the pale in difference from paradigm exemplars of F , given the subject's point of view and other relevant circumstances of perception. Of course this would secure the right result, but only by brute force and without any kind of explanation of the limits on the nature and extent of any errors involved in perceptual presentation. The notion of being 'beyond the pale' is simply manufactured to respond to these limits without giving any account of them or why they obtain.

What this suggests to me at least, and what I endorse wholeheartedly in my own (OV) in ch. 5 below, is that the question of which mind-independent physical objects are presented in perceptual experience is prior to and more fundamental than any question of how those things may look from the relevant point of view and in the circumstances in question. Perceptual presentation is absolutely not to be reduced to or manufactured out of supposedly fundamental perceptual contents of this kind. Given that a particular object is indeed presented in perception. Limits on the ways it may look intelligibly follow from its nature and the subject's point of view and various other relevant circumstances of perception.

I turn now to singular perceptual contents, where my conclusion is exactly the same. Once again, there is no restriction whatsoever of the kind we are interested in, so far as I can see, on the nature or extent of any error involved in a content of the form

<Fa>, at least insofar as this is conceived as the content of thought.³² Suppose we name the figure on a diagram you have in your pocket that happens to be of the (ML) lines 'A1'. 'A1 consists of two lines of unequal length' and 'A1 consists of a perfect circle' are perfectly coherent thought contents. You may even present me with two pieces of paper one inscribed with each of the sentences expressing these two contents, and ask me to choose between them. I will be wrong whichever I choose. For the lines are in fact equal in length. Indeed I can see no significant difference between them that is adequate as things stand to account for the fact that only the former is a candidate content for a genuine perceptual illusion that presents the (ML) lines in experience. Now, recall the account of perceptual presentation given above for singular contents of this kind: <Fa> presents a just if a is the reference of the singular content component 'a'. So the singularist proponent of (CV) may attempt to replicate the fact that only <a consists of two lines of unequal length> is a candidate content for a genuine illusion presenting the (ML) lines in experience by insisting that the singular components of perceptual contents, unlike some thought contents at least, are demonstratives whose successful reference to particular mind-independent physical objects places limits upon the nature and extent of any errors involved. Of course this would secure the right result, but it depends entirely upon the claim that the relevant demonstratives refer only in the absence of errors beyond the limits in question. This seems to me absolutely right; but it is a consequence of two further facts: first, that the relevant demonstratives are perceptual demonstratives that refer to the objects presented in perception; and second, that there are limits on the nature and extent of any errors involved in cases of genuine perceptual presentation. Thus, far

³² Perhaps there are categorical constraints. See Magidor (forthcoming) for very helpful sceptical discussion. In any case, such categorical constraints would clearly not suffice to secure the relevant data concerning perceptual presentation.

The Content View

from explaining the phenomenon in question, of limits on any error in perceptual presentation, this account simply presupposes and recycles it.

Once again then, it seems to me that the presentation of mind-independent physical objects in perception is a more fundamental phenomenon than the specific ways such things may look on any particular occasion, and is not itself to be explained, as (CV) is compelled to explain it, in terms of the representational contents of perceptual experience. The ways physical objects look is a consequence of, and thereby constrained by, the more basic phenomenon of the presentation of particular physical objects in perception. Perceptual presentation provides the ground for the very possibility of particular mind-independent physical objects looking any way at all. This is the source of the limits on the nature and extent of any errors involved in genuine perceptual illusion. Hallucination is to be given a quite different account. These are all topics for ch. 5. In § 4.3 below I discuss a second problematic feature of (CV): the essential generality involved in its commitment to predication at the most fundamental level of our perceptual relation with the physical world.

4.3 The Generality of Predication

(CV) characterizes perceptual experience by its representational content. In doing so, it retains even at the level of our most basic perceptual encounter with the physical world around us a key feature of (IM), namely, that content admits the possibility of falsity: the world might not actually be the way that a given content represents it as being; and the scope for such error may in certain cases be quite drastic. It is often assumed to be a major benefit of (CV) that this feature may be put to use in its

The Content View

explanation of perceptual illusion and hallucination. I argued in § 4.2 above, though, that this assumption is mistaken, and that the possibility of falsity at the most fundamental level of our perceptual relation with the physical world is a net cost, not a benefit, to (CV): it stands in the way of an adequate account of the genuine presentation of physical objects in perception and illusion, as against hallucination.

The current section proceeds as follows. First, I explain the way in which (CV) is committed to a certain generality in its most fundamental account of our perceptual relation with the physical world. Second, I argue that this generality is in tension with the datum that particular mind-independent physical themselves are presented to us in perception.

According to (CV), the most fundamental material for understanding our perceptual relation with the physical world comes from the worldly representational content of perceptual experience: the way such experience represents things as being in the world around the subject. The generality of the predication involved in perceptual content is most obvious in connection with (IM): the thought content that a is F. Here a particular object, a, is thought to be a specific general way, F, that such objects may be, and that infinitely many qualitatively distinct possible objects are. 'F' is associated with a specific general condition; and the particular object, a, is thought to meet that very condition. I claim that this generality of predication that is explicitly present in the simple thought content that constitutes (IM) is essential to the truth-evaluability of content in general, which provides the key motivation for the (CV) account of illusion and hallucination. The 'particulars' involved need not necessarily be persisting material objects, or, indeed, 'objects' of any kind. Even the most abstract formulation

of a truth-evaluable content as that things (or the relevant realm of reality) are (is) thus and so (as opposed to some other way), displays the particular/general combination.³³ Thus, I claim that the representation of some item or items as being some more or less specific general way is a crucial characteristic of the contents that are fundamental to (CV). Furthermore, the key feature of such contents for my purposes is the generality of their predicational component: this places a demand upon items of the relevant kind that may be satisfied equally by a whole range of numerically and qualitatively distinct individuals of that kind. Of course, if ‘x is F’ is the predicate in question, then there is a way in which the items meeting the associated demand may not qualitatively differ, namely in all being F; but there are (possibly indefinitely) many other ways in which they may be qualitatively distinct. I claim that this generality of predication, which is crucial to (CV) perceptual contents, precludes an adequate account of the presentation of particular mind-independent physical objects in perception.

Suppose that you see a particular red ball – call it ‘Ball’. According to (CV), your perceptual experience of Ball is to be characterized by its representational content. Let us assume that this content makes singular reference to Ball.³⁴ Your experience therefore represents that Ball is a specific general way, F, that such objects may be. Whichever way this is supposed to be, its identification requires making a determinate specification of one among indefinitely many possible generalizations from Ball itself. Ball has colour, shape, size, weight, age, cost, and so on. So perception must

³³ Even this very general claim has its opponents. For a helpful and historically informed survey of the debate beginning with Ramsey (1990b), see Macbride (2005).

³⁴ So far as I can see this assumption may be made without loss of generality, since the argument would be unaffected if the content in question were quantificational instead, instantiated in some way by Ball.

The Content View

begin by making a selection amongst all of these, according to (CV). Furthermore, and far more importantly for my present purposes, on any given such dimension – colour, or shape, say – the specification in experience of a determinate general way that your perception supposedly represents Ball as being requires further crucial abstraction. Supposing that your experience is veridical, it must be determinate to what extent, and in which ways, Ball’s actual colour or shape might vary consistently with the truth of the relevant perceptual content. This is really just to highlight the fact that (CV) is committed to the idea that your perceptual experience has specific truth conditions, which go beyond anything fixed uniquely by the actual nature of the particular red football – Ball – that you see.

According to (CV), then, perception – even perfectly veridical perception, whatever exactly this may be – does not consist in the simple presentation to a subject of various constituents of the physical world themselves. Instead, it offers a determinate specification of the general ways such constituents are represented as being in experience: ways that other numerically distinct such constituents, qualitatively distinct from those actually perceived by any arbitrary extent within the given specified ranges, might equally correctly – that is, truly – be represented as being. Any and all such possible alternatives are entirely on a par in this respect with the object supposedly perceived, so far as the way things are presented in experience according to (CV) is concerned. In regimenting our most fundamental perceptual relation with the physical world by appeal to an abstract act of predicational classification or categorization in this way, (CV) therefore fails to provide any adequate account of the direct perceptual presentation of particular mind-independent physical objects themselves in perceptual experience. For there is no genuine

The Content View

substance to the idea that it is the actual physical objects before her that are presented in a person's perception, rather than any of the equally truth-conducive alternative possible surrogates. She may supposedly be referring to a privileged such entity in thought, but it is hard to see how it is evidently that thing, rather than any other, that is truly presented to her in experience. The generality of the predication that essentially occurs in the complete contents central to the (CV) account of our fundamental perceptual relation with the physical world bars any adequate account of the way in which perception presents particular mind-independent physical objects to us in experience.

An initial reply to this line of objection would be to protest that (CV) is perfectly consistent with the idea that the actual constituents of the mind-independent physical world are presented to us in experience. For the constituents in question are the very properties of physical objects that are ascribed to them by the predicational components of complete perceptual contents.³⁵ I have two points to make in response. First, this reply as it stands is an explicit admission of the failure of (CV) to account for the fundamental phenomenon that was agreed on all sides from the outset of the presentation in perceptual experience of particular mind-independent physical objects. Second, my own nominalist view is that the general properties of physical objects are not genuine constituents of the mind-independent world itself. Of course physical objects have all the properties that they have; and their having them is in many cases a perfectly mind-independent truth, in no way constituted by whether or not they seem to have them to us or to anyone else. Still, it is a mistake to postulate constituents of reality corresponding to the predicates used in expressing such truths on a par with the

³⁵ Thanks to Tim Crane for this suggestion.

objects themselves corresponding to the relevant singular terms and of which the predications in question are truly made.³⁶ Hence the current reply to my objection in my view threatens even the more general idea that perception involves the experiential presentation of genuine constituents of the mind-independent physical world itself.

A second reply on behalf of (CV) seeks to undermine the particular conception of the generality of predication upon which my objection is based.³⁷ I claim that this imposes a demand upon the object perceived that might equally be satisfied by any number of numerically and qualitatively distinct alternatives, and that this obstructs a satisfactory account of the sense in which that very object itself is presented in experience: that thing and nothing else is actually seen. Suppose for simplicity of presentation again that the content in question is that a is E, and that a is in fact E. The objection depends upon the fact that ‘x is E’ may equally be satisfied by a whole range of alternatives to a that differ from a not only numerically but also, and in (indefinitely) many ways, qualitatively too. The reply is that the range of alternatives to a do not differ from a qualitatively at least, since the relevant way that they all are is E: they are all precisely “the way ... [that a] visibly is” (McDowell, 2008a, p. 204). This is just the point I made above, though, that there is indeed one respect in which all and only the objects satisfying ‘x is E’ may not qualitatively differ: they are E; but this is perfectly consistent with the equally evident fact that there are (indefinitely) many other ways

³⁶ For an excellent overview of the issues here and very helpful bibliographical references, see Rodriguez-Pereyra (2008). Nominalism is of course a massive issue and I cannot possibly engage with it adequately here. I mention my own commitments for completeness; the dialectical weight at this point is adequately borne by my first response.

³⁷ This objection is due to McDowell (2008a, § 4).

The Content View

in which the items satisfying any such predicate may be qualitatively distinct; and this is all that my objection requires.

A third reply to the current objection to (CV) is that the idea that it is the actual physical objects before her that are presented in a person's perception, rather than any of the equally truth-conducive alternative possible surrogates, is to be substantiated by the fact that it is those very things, rather than any alternative such surrogates, that play an appropriate role in causing the perceptual experiences in question. I argued in § 4.2 above that it is a necessary condition upon the subjective presentation of a particular mind-independent physical object, o, in perceptual experience that there should be some degree of match between the way things look in such experience and that very object, o, itself, although this must of course also be compatible with the possibility of various kinds of illusion. The current proposal is that this must be supplemented with a causal condition according to which o itself is suitably implicated in the production of an appropriately matching experiential representation if that very thing is genuinely to be presented in perception. This is what it is for o to be so presented. The result is perhaps the most natural development of the (CV) approach to perception.

An initial counter would be to query whether this is the right kind of response to address the objection raised. The worry was that, according to (CV), perceptual experience itself appears to be subjectively silent on whether o itself is determinately presented in experience in the face of the equal claim of a whole range of appropriately matching qualitatively distinct alternatives. The current response simply acknowledges this subjective silence and hopes to make up for it by appeal to the

The Content View

causal history of the actual experience involved. The intuition motivating the objector, though, is that perceptual presentation of a particular mind-independent physical object is precisely a subjective matter: it is as central to the way things are experientially for the subject of perception as anything could be that he is presented with a particular object before him. This initial counter is relatively inconclusive as it stands, though. For the whole point of the present (CV) response is explicitly to insist that the phenomenon of perceptual presentation should on reflection be reconstrued as a composite of relatively independent subjective-experiential and causal components. At this stage, then, there is danger of a static stand-off between opponent and proponent of the (CV) approach.

I wish to argue further that the proposed decomposition is in itself seriously problematic. Hence the initial intuition that perceptual presentation is absolutely essential to the intrinsic subjective nature of perceptual experience is absolutely sound and the current objection to (CV) therefore stands. Appeal to the causal role of various mind-independent physical objects in the production of perceptual representations is incapable of saving the core phenomenon of the experiential presentation of particular such things in perception in the context of (CV). The argument is a direct application of a very powerful general strategy to undermine the decomposition or ‘factorization’ of cognitive phenomena that is due to Timothy Williamson (1998, 2000, forthcoming).

Suppose that being presented in perception with a particular mind-independent physical object, o, a large white cube, say, is a composite condition, with an internal experiential matching component, I, and an external environmental-cum-causal

component, E. Thus being presented with \underline{o} is equivalent to I&E. Now, on the present version of (CV), a person may certainly be presented with \underline{o} by having an experience in which there looks to be a large white cube immediately to her right that is caused by \underline{o} being immediately to her right in ideal viewing conditions. So having an experience caused by \underline{o} being immediately to one's right in ideal viewing conditions is compatible with I&E and hence with E. Call this environmental-cum-causal condition – having an experience caused by \underline{o} being immediately to one's right in ideal viewing conditions – E'. Similarly, a person may also be presented with \underline{o} according to the present version of (CV) by having an experience in which there looks to be a small grey dot a long way off to her left that is caused by \underline{o} being a long way off to her left in poor lighting conditions. So having an experience in which there looks to be a small grey dot a long way off to one's left is also compatible with I&E and hence with I. Call this internal subjective experiential condition – having an experience in which there looks to be a small grey dot a long way off to one's left – I'. Next, assume that any possible internal experiential condition is compatible with any possible environmental-cum-causal condition.³⁸ It follows that I&I' is compatible with E&E'. Since this is a condition in which I&E, it is a condition in which \underline{o} is presented in perception. Yet it is a condition in which a person has an experience in which there looks to be a small grey dot a long way off to her left that is caused by \underline{o} (a large white cube) being immediately to her right in ideal viewing conditions; and this I claim is not compatible with being genuinely subjectively presented with that very cube, \underline{o} , in experience, even given the acknowledgement above of a certain

³⁸ This assumption may be questioned. It is a very natural assumption to make on the present version of (CV), according to which the phenomenon of being perceptually presented with a particular mind-independent physical object is to be construed as a composite of relatively independent subjective-experiential and causal components. I leave it as an explicit assumption of my argument here against this (CV) proposal.

degree of contingency in the limits upon the nature and extent of the errors in perceptual appearance that are compatible with genuine presentation in experience. It therefore follows that being presented with o itself in perception cannot be decomposed into any pair of logically independent internal experiential and environmental-cum causal components.³⁹

Thus the strategy of giving substance to the idea that it is the actual physical objects before her that are presented in a person's perception, rather than any of the equally truth-conducive alternative possible surrogates by appeal to the fact that it is those very things, rather than any alternative such surrogates, that play a role in causing the perceptual experiences in question cannot be sustained. The present objection to (CV) stands. In regimenting our most fundamental perceptual relation with the physical world by appeal to an abstract act of predicational classification or categorization, (CV) fails to provide any adequate account of the direct perceptual presentation of particular mind-independent physical objects themselves in perceptual experience.

However automatic or natural the general classification essentially involved in the perceptual predication central to (CV) may be, it still constitutes an intrusion of conceptual thought about the world presented in perception into the account of our fundamental perceptual relation with the physical world. According to my own positive position (OV) as I develop this in ch. 5 below, the selective categorization of particular constituents of physical reality enters the picture of a person's relation with the world around her only when questions of their various similarities with, and differences from, other such things somehow become salient. This categorization

³⁹ See Williamson (1998 and forthcoming) for detailed discussion and application of the anti-decomposition strategy also implemented here.

The Content View

depends upon a more fundamental level of the presentation to her of particular mind-independent physical objects in perceptual experience rather than constituting an essential part of that very perceptual presentation. Perceptual presentation itself constitutes the fundamental ground for the very possibility of any such general classification of the constituents of the physical world that are in this way made available in perceptual experience. This most fundamental perceptual phenomenon is not to be characterized in terms of any kind of predicative generality, as it is according to (CV). It is in itself wholly particular in my view.

Proponents of (CV) may hope to soften the impression that their characterization of perceptual experience by its representational content in this way constitutes a mistaken importation of selective abstraction and categorization into the most fundamental account of our perceptual relation with the physical world, along the following lines. Genuine – that is veridical – perception presents a person with various constituents of the physical world themselves: particular mind-independent physical objects. Still, it must be acknowledged that this always involves less than perfect acuity. There is a determinate range of respects in which those very things might have been different without any relevant difference in the impact made by them upon the subject in question. Thus her perception is bound to involve a degree of generality. The general way that her experience represents such things as being is precisely that way that would determine the resultant perceptual content as true if and only if the relevant worldly constituents were as they actually are in every respect, or were instead different in any of the specific respects within this range.⁴⁰

⁴⁰ I think that McDowell's (2008a, § 4) objection to this disjunctive formulation rests on a confusion. Of course, if it is already given that my experience represents a as E, then the second disjunct of a disjunctive statement of its truth conditions is clearly

Such hope is in my view misguided. For this proposal faces a number of serious difficulties.⁴¹

First, and most importantly, it has more than a whiff of circularity. The suggestion is that perceptual experience is to be characterized by its representational content, which is in turn to be identified by a certain procedure that takes as its starting point a worldly situation in which that very content is supposed to be determined as true. That is, the truth conditions definitive of the experiential content in question are to be specified by a kind of generalization from a paradigm instance of its actual truth. Yet how is it supposed to be determined what is to count as such an instance of its truth, for the purposes of generalizing to these truth conditions, in advance of any specification of those very conditions? This proposed procedure for the

otiose: either a is F or it is any way that would be indiscriminable from being F by the question whether or not it is F. The point of the current proposal on behalf of (CV), though, is supposed to be to motivate, or at least defend, the position by offering an approach to determining truth conditions for perceptual contents, or by providing a constitutive account of what the truth conditions of a given perceptual experience are. The basic idea is that in any case in which a person has a particular object successfully in view, then the predicational component of the content of the relevant experience imposes a determinate demand upon objects in the world in general that they be qualitatively exactly like the particular object in question, or sufficiently similar not to make a relevant difference to the experience along the property dimension in question, such as colour or shape. This fixes the ‘band-width’, as it were, of the perceptual experiential predication along that dimension. As I argue below, the proposal is unacceptable, and for reasons long the lines McDowell suggests; but this is not to the benefit of (CV), but entirely to its detriment.

⁴¹ These are not simply technical or empirical issues of detail that may be regarded by proponents of (CV) as either relatively independent of the argument sketched above from imperfect perceptual acuity or pertaining to a challenge that need not be taken on to specify exactly what the content of any given perceptual experience is. They are intended rather as objections in principle to the proposed theoretical regimentation of genuine phenomena of imperfect acuity in terms of the generality of perceptual representational content.

The Content View

characterization of perceptual experience cannot even get started unless it has already been completed. It is therefore either useless or unnecessary.

Second, suppose that we have somehow determined that the case before us is one of genuine – that is, veridical – perception, rather than illusion or hallucination; and suppose, further, that we have some way of fixing the actual constituents of the subject's environment that are experientially accessible to her. The proposed specification of the representational content of her experience then proceeds as follows. Its truth conditions are satisfied if and only if, things are precisely as they actually are, or they are different in any of the various respects in which they might have been different without making any relevant difference to their impact upon her. This immediately raises the question which differences are relevant, in the impact made upon the subject. Any change in the worldly constituents in question makes a difference of some kind, even if this is only characterized in term of her relational embedding in a different environment in which that change obtains. Relevant changes of the environment, though, are those that transform the world from a condition in which the initial target content of her perceptual experience is to be regarded as true, to one in which it is to be regarded as false. So the question of which worldly differences are relevant is clearly crucial. I cannot establish here that no satisfactory account of what makes such differences relevant can possibly be given. So this line of argument is bound to remain a challenge to the present defence of the way in which (CV) imports generality into the characterization of our most fundamental perceptual relation with the physical world, rather than a conclusive refutation. Still, the following four proposals are clearly problematic.

The Content View

1. A worldly change is relevant iff it makes an intrinsic physical difference to the subject's perceptual system. This is plausibly neither necessary nor sufficient for the world to change its condition from one in which the subject's initial perceptual content is true to one in which it is false, according to (CV). Any trace of that form of externalism in the contents countenanced for perceptual experiences on which these fail to supervene upon a subject's intrinsic physical condition simply consists in the denial of its necessity; and some such externalism is widely endorsed by proponents of (CV) (see e.g. Pettit and McDowell, 1986; Burge, 1986; Peacocke, 1992, ch. 3; and Davies, 1997). On the other hand, the idea that an effect on the intrinsic physical condition of the subject's eyes, say, is sufficient to transform any worldly condition in which a given experiential content is veridical, into one in which it is not, surely individuates perceptual contents far too finely. For we are notoriously capable, from a very early age, of representing crucial environmental constancies, such as shape and colour, as such, across variations often far more significant than these; and there is a vast amount of information that is picked up by the visual system, for example in unconscious perceptual priming, that we are not conscious of at the time, but that may be brought to light by appropriate later testing.⁴² The required (CV) response that the overall perceptual content changes in some way in every such case strikes me as rather desperate.

2. A worldly change is relevant iff it actually makes a difference to the way the subject believes things are out there. Again, this is arguably neither necessary nor sufficient for a worldly change to be relevant in the required sense. If she is suitably preoccupied with the colour of an object before her, for example, variation in its

⁴² See Forster and Davis (1984) for representative examples and discussion.

shape, say, to an extent that would render her current perceptual representation of this shape false, may nevertheless make no impact whatsoever on her actual beliefs about it. On the other hand, (CV) must presumably allow for the possibility, at least, that a change in the way things are in the world around her makes a difference to the subject's beliefs about the world entirely independently of the way things are actually represented as being in her experience. Indeed, proponents of the present version of (CV) have no alternative that I can see but to appeal to this very idea, of worldly changes affecting a person's beliefs otherwise than by influencing the content of her experiential representations, in explanation of the systematic effects of various masked stimuli, for example.⁴³

3. A worldly change is relevant iff it actually does make, or might, without modifying its intrinsic physical effects upon the subject, have made, a difference to the way she believes things are out there. Perhaps a possible effect upon the subject's worldly beliefs of this kind is a necessary condition of any worldly change which renders a previously veridical experiential content false, although any such possibility is intuitively causally explanatorily grounded in experiential change rather than constitutively explanatorily of it. Still, since the current condition upon the relevance of a world difference is strictly weaker than the previous one that I argued is insufficient, it must be insufficient too: rapidly masked stimuli may actually (hence actually-or-possibly) affect a subject's beliefs about the world without showing up in any way in experience.

⁴³ These are cases in which very briefly exposed items systematically affect the beliefs of subjects who deny any conscious perception of them. See Breitmayer and Ogmen (2006, esp. ch. 8) for a comprehensive survey and bibliography.

4. The nature of this insufficiency suggests a fourth approach that is surely in the vicinity of what (CV) needs here, although I shall argue that it is either circular or independently objectionable in the current context for reasons that we have already seen. The proposal is that a worldly change is relevant, in the required sense, iff it makes a difference to the subject's experience of the world. This immediately raises the question, though, how such differences in experience are to be characterized. I can see just two possibilities, neither of which is acceptable.

First, they are differences in its representational content. The idea would presumably be something like this. A person has a perceptual experience, and we are presuming for the sake of the argument both that it is veridical and that we have identified the worldly objects and their features that it concerns. In order to determine its specific representational content, we are to consider the various ways in which these very objects might have been different with respect to such features. The content will be true in all of those cases in which such variation does not change its content. In other words, in order to carry out this procedure for the determination of perceptual content, we have already to have fixed that very content. So the procedure is clearly unacceptably circular. One might insist in reply to this simple objection of circularity that there are distinct levels of content in perception. Thus, 'higher' levels may be determined in the way outlined above by appeal to the various worldly changes that leave the 'lower' levels unaffected. This hierarchy of levels presumably has a lowest member in every case, though, and the objection of circularity applies directly at that level, and therefore applies indirectly throughout. So, again, this first implementation of the idea that worldly changes are relevant in the required sense iff they make a difference to the subject's experience of the world is unacceptable.

Second, the differences in experience, by reference to which the required notion of relevant worldly variation is to be characterized, are differences in its intrinsic character, where this is to be elucidated prior to and independently of any question of its representational content. This is explicitly inconsistent with (CV) as I define it, though. For the whole point of the current proposal is to avoid the circularity of the previous suggestion by appeal to a characterization of perceptual experience more fundamental than its representational content, and on the basis of which such content may be understood. Yet the definitive commitment of (CV) is the claim that the most fundamental account of our perceptual relation with the physical world is to be given directly in terms of the complete representational contents of perceptual experience. So this second implementation of the idea that worldly changes are relevant in the required sense iff they makes a difference to the subject's experience of the world is equally unacceptable in the current context.

The basic idea here is much more conducive to the indirect realist setting in which perceptual experience consists most fundamentally in acquaintance with certain mind-dependent direct objects. Such experience nevertheless represents the way things are in the mind-independent physical world around the perceiver. In order to determine its characteristic representational content in any given case, we begin with a situation in which the perceiver has an experience with the direct objects in question and everything is operating normally, however exactly normal operation is initially to be identified. We then consider gradual changes in the world around the perceiver along various property dimensions up to the point at which the relevant direct objects of experience themselves change as a result. The truth-conditions of the perceptual

The Content View

content in question are that the world be within that range around its initial condition along each of the property dimensions in question. The result is a form of indirect realism according to which the mind-dependent direct objects of experience are natural signs of various mind-independent worldly states of affairs (Ayers, 1993, vol. I, ch. 7). The extent of the generality introduced into perceptual content corresponds to the degree of acuity in the signing system. We saw in ch. 3 above that any such position is not only inconsistent with (CV) but also untenable in its own terms.

Still, there may be something important and right about the basic idea here that the fundamental nature of our perceptual experience of mind-independent physical objects is the prior and independent ground of the way that such objects actually do and possibly might look to us in perception. My aim in chs. 5-7 below is to elucidate an account along precisely these lines in some detail.

Unfortunately for proponents of (CV), though, the basic insight is circular if it attempts to combine this with a characterization of experience itself exclusively in terms of its representational content; and it collapses into an untenable indirect realism if it attempts to supplement this content-characterization with any appeal to more basic, mind-dependent direct objects of experience. The right response to this impasse, in my view, is to reject (CV) altogether. The fundamental nature of perceptual experience does indeed provide the explanatory ground of the ways things do and might look to subjects from various points of view and in various circumstances of perception. It does so, though, not by serving up fully formed representational contents, but, rather, by presenting her directly with the actual

The Content View

constituents of the physical world: persisting mind-independent physical objects themselves.

Of course there are many more possible proposals than the four that I have considered here for the way in which the most basic representational contents of perceptual experience are supposed to be determined according to (CV). Still, I do think that one might perfectly reasonably conclude from this representative sample of failures that the current attempt to defend (CV)'s commitment to predicational generality faces a very serious challenge. The presence of predicational components is essential to the representational contents that are definitive of the (CV) account of the fundamental nature of our perceptual relation with the physical world. The challenge that I offer to its proponents is to explain at least in principle how to arrive at a specification of the various worldly changes that are relevant to the transition from truth to falsity in any given perceptual experience. For these are crucial to determining the significance of the general predicational components involved in its content.

In any case, all of these problems for giving a specific account of how the most basic representational contents of perception are supposed to be determined according to (CV) combine with the objection of principle that I set out earlier in the present section. The generality of predication that is crucial to the representational contents that (CV) regards as fundamental to our perceptual relation with the physical world around us is in fact incompatible with any adequate account of the presentation of particular mind-independent physical objects to us in perceptual experience. I argued in § 4.2 that the possibility of falsehood in perceptual contents that many regard as a primary motivation for (CV) also stands in the way of any proper acknowledgement

The Content View

of fact that mind-independent physical objects are presented to us in perception.

These tensions came out further in my discussion in § 4.1 of the particular difficulties faced by any attempt on behalf of (CV) to exploit the possibility of false perceptual contents in accounting for illusion and hallucination. I therefore conclude that (CV) has to be rejected. In the three remaining chapters and in the light of all that has gone before, I offer, elucidate and defend my own positive account of our most fundamental perceptual relation with the physical world.