

SEEING AND EXPERIENCING: THE REVELATION OF
PARTICULARS IN VISUAL PERCEPTION

by

Alexander B. Hughes

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Philosophy)
in The University of Michigan
2007

Doctoral Committee

Professor Peter Ludlow, Co-chair
Associate Professor Eric P. Lormand, Co-chair
Professor Robert G. Pachella
Professor Kendall Walton

ACKNOWLEDGEMENTS

Thanks are due to my committee for their tireless efforts. But for their ministrations, this dissertation would not exist. Eric Lormand deserves special recognition for his help. Since the beginning of this project, Eric has been an invaluable source of criticism.

Though he does not know it, Karson Kovakovich deserves thanks for the inspiration and education I obtained from our weekly meetings in the winter of 2003. It was the worst winter of my life, but I learned more in conversation with him than from months of solitary scholarship.

I am, of course, indebted to my friends, family, and loved ones for their patience and support. For her support during the formulation of the project, some debt is owed to Jessica Rett. For their encouraging advice, Remy and Lori should be thanked. For their unstinting optimism, my parents deserve recognition. For all the labor they performed on my behalf, thanks are due to my past selves. Such selflessness! Finally, for her patient forbearance, constant companionship, and sage advice during the last few years, Liz Goodnick deserves the greatest share of my gratitude. Thank you Liz.

CONTENTS

ACKNOWLEDGEMENTS	ii
CHAPTER 1: GETTING INTO THE ISSUES: SEEING AND EXPERIENCING	1
1.0: INTRODUCTION	1
2.0: THE STANDARD VIEW, EXTERNALISM, AND VISUAL CONTENTS	6
3.0: THE INSURGENT VIEW: DISJUNCTIVISM	16
CHAPTER 2: PERCEPTUAL DEPENDENCE: FINDING A PLACE FOR COUNTERFACTS	21
PART 1: COUNTERFACTUAL DEPENDENCE AND SEEING	22
1.1: COUNTERFACTUAL DEPENDENCE AND SEEING	22
1.2: THE CENSOR CASES: <i>CONTRA</i> LEWIS	28
1.3: DO YOU SEE IN CENSOR CASES? LEWIS' ASSAULT ON COMMON SENSE	29
1.4: DO YOU SEE IN CENSOR CASES? AN ANALOGY FROM ACTION THEORY	31
1.5: THE REVERSE CENSOR, THE RAY GUN	35
1.6: MY VIEW	36
1.7: BRUCE LECATT: STEPWISE DEPENDENCE	39
1.8: THE SUPER CENSOR: <i>CONTRA</i> LECATT	42
PART 2: PRINCIPLED OBJECTIONS	44
2.1: MCLAUGHLIN'S WORRIES	44
2.2: ASSESSING THE MCLAUGHLIN-STYLE OBJECTIONS	47
PART 3: PHOTOGRAPHS AND ROBOTS	54
3.1: THE CASE OF TIM AND TOM	54

3.2:	A CONCESSIVE RESPONSE: TYE’S PROPOSAL	55
3.3:	PHOTOGRAPHIC REALISM AND SEEING TIM	58
3.4:	PHOTOGRAPHIC REALISM	59
3.5:	SEEING TIM	64
PART 4:	THE THEORETICAL POSSIBILITY OF ELIMINATING REFERENCE TO CAUSATION	66
PART 5:	THE GENERAL GENERALITY PROBLEM AND ME	68
	CONCLUDING REMARKS	80
	CHAPTER 3: THE POSSIBILITY OF VERIDICAL HALLUCINATION	82
0.1:	INTRODUCTORY COMMENTS	82
0.2:	THE ARGUMENT FROM VERIDICAL HALLUCINATION	83
PART 1:	SOTERIOU’S CHALLENGE	85
1.1:	FROM VERIDICAL HALLUCINATION TO VERIDICAL MISPERCEPTION	85
1.2:	ASSESSING SOTERIOU’S ARGUMENT	92
1.3:	THE DEFENSE OF VERIDICAL HALLUCINATION AND THE ARGUMENT FROM COMPETING OBJECTS	101
1.4:	SHIFTING CONTENTS: REPLACEMENT AND ADDITION	104
1.5:	THE POSSIBILITY THAT INVOKING CAUSAL DEPENDENCE IS GRATUITOUS	106
1.6:	CONCLUDING REMARKS	110
PART 2:	SEARLE AND VERIDICAL HALLUCINATION: ARE SEEING RELATIONS REPRESENTED IN THE CONTENTS OF VISUAL EXPERIENCES?	111
2.1:	SEARLE’S ACCOUNT	112
2.2:	ASSESSING SEARLE	115
2.2.1:	ASSESSING SEARLE: CONTENTS, COUNTERPARTS, AND PHENOMENOLOGY	118
2.2.2:	ASSESSING SEARLE: VISUAL CONTENTS AND THE MENTAL ENVIRONMENT	135
2.2.2.1:	THE MENTAL ENVIRONMENT: DIFFERENCES AMONG ATTITUDES/SENSORY MODALITIES	136

2.2.2.2:	THE MENTAL ENVIRONMENT: PERCEPTION, THE PROPOSITIONAL ATTITUDES, AND MENTAL ANAPHORA	139
	CONCLUDING REMARKS	152
	BIBLIOGRAPHY	154

CHAPTER 1

GETTING INTO THE ISSUES: SEEING AND EXPERIENCING

1.0: INTRODUCTION

Judgment – when its exercise rises to the happy status of knowledge – is clearly one of our chief modes of access to reality. Accordingly, it is not a surprise that epistemology is among the perennial concerns of philosophers. To the extent that access to reality is a good thing, it is a good thing to reach an understanding of it. This is because understanding might improve our capacity for judgment, but also because it is worth understanding what is valuable. Thus it is reasonable for us to inquire into the psychology, metaphysics, and value theory of knowing judgment.

This dissertation, however, is focused on a distinct mode of access to reality – the perceptual mode of access. The value of knowledge (and hence the value of understanding it) rests on the fact that it is a mode of access to reality. Similarly, since perceptual modes of access to reality are modes of access to reality, it is valuable to understand perception.

How, then, is it that reality is revealed to us when we cast our gaze about? How is it that the region of the cosmos (and its features) around us is disclosed in our perceptual experiences? Taken one way, these are scientific questions, to be determined by examining what happens when we see. However, I mean these to be philosophical questions: I want to know what constitutes whatever the scientists discover as the perceptual disclosure of reality.

Our sensory awareness acquaints us with ever so many different aspects of reality. For some examples, one's sensory awareness can constitute acquaintance with particulars, properties, and events. I propose, however, to focus on the way in which particulars are revealed in experience. Part of the reason for this is purely autobiographical: this question – what is it to see a thing – is how I began thinking about these issues. Additionally, I suspect that the revelation of particulars is explanatorily prior to the revelation of, for example, events – though perhaps not explanatorily prior to the revelation of properties.

At any rate, some reason for focusing on the revelation of particulars derives from its intrinsic interest. Moreover, perception generally – and visual perception in particular – appears to play a crucial role in fixing the contents of our mental states. Imagine a creature who is equipped with the capacity to detect features, but without the capacity to think about particulars. Such a creature would probably think many true things, but would remain out of touch in a peculiar way. To the extent that it matters that a creature is in touch with one thing rather than another, we should feel a kind of pity for it. Our romantic love, for example, aspires to bear on one person – and not, for example, any one of that person's possible twins. One plausible partial answer to the question of what constitutes our attitudes as bearing on this or that particular is this: we must first perceive a thing, if our attitudes are to bear on it. Of course, this answer needs refinement, but if something like it is right, then perception is the beginning of our dealings with particular things.

As I will argue in Chapter 3, the fact that my experience purports to reveal, for example, the texture of your hair can depend on its being you I am seeing right now. If I

were seeing your twin, then the content of my visual experience – that which constitutes its visually seeming to me things are thus-and-so – would be such that the state of your twin’s hair would determine whether my experience is veridical. Thus the perceptual facts can make a difference to the contents of visual experience – to which particulars are required to possess such-and-such features, if one’s experience is to be veridical.

How, then, is it that particular things are revealed to us when we cast our gaze about? What constitutes our seeing one thing, rather than another or nothing at all (as in the case of hallucination)? What constitutes an experience as a perception of, say, Fred rather than his mother, Zenobia? What constitutes an experience as a hallucination, rather than a perception of some particular?

Grice's thought experiments motivate the view that an initial answer to this question requires that a perceptual experience causally depend on its objects. (Grice 1961) But, the Gricean dialectic continues, as cases of causal deviance show, causation is not enough. In order to constitute a perception of a thing, one's experience needs to be caused by it in the right kind of way. Grice leaves it open what the right kind of causal dependence might be. Progress in formulating a dependence condition on perception consists in saying more about what the right kind of dependence is. The hope is that we can tell a story about what feature every perceived particular possesses and every unperceived particular lacks. This story, moreover, will be the story of what constitutes the revelation of things in our experience.

This dissertation is organized around engaging with the Gricean arguments. Grice leaves it open what the right kind of causal dependence might be. The business of Chapter 2 is to say more about what the right kind of causal dependence is. In particular,

I favor a version of counterfactual dependence theory. The core commitment of a counterfactual dependence theory of visual perception is that seeing a thing is constituted by its being true that one's experience tracks the state of the thing across nearby modal space. Thus, for example, if you are seeing a red ball, then were the ball blue (instead of red), then your experience would represent it as red. However, counterfactual accounts suffer from two kinds of problem. First, they suffer from counterexamples involving counterfactual interventions – interventions that do not happen, but would, were the actual situation different. Second, they suffer from principled objections. That is, it has been argued by Brian McLaughlin and Colin McGinn that there are reasons to doubt that any counterfactual account can succeed. I address both of these problems. First, I defend the counterfactual theory against the counterexamples by proposing a process-relative counterfactual account – were the seen thing different (in appropriate ways) and were the process actually eventuating in your experience operational, then your experience would represent those differences. Second, I argue that we need not be convinced by McLaughlin and McGinn's criticisms of counterfactual theory.

In his article, 'The Causal Theory of Perception', Grice offers two arguments for the claim that the seeing relation is a kind of causal-explanatory relation. (Grice 1961) I call them the *Argument from Veridical Hallucination* and the *Argument from Competing Objects*. Chapter 3 engages with the Gricean arguments by defending the soundness of these against two sources of doubt about the possibility of veridical hallucination. The first half of Chapter 3 defends the Gricean arguments against Matthew Soteriou's argument that veridical hallucination is an impossibility.¹ On his account, if you accept the possibility of veridical hallucination, then you must accept that an experience could

¹ The argument is contained in his article, 'The Particularity of Visual Perception'. (Soteriou 2000)

be both veridical and non-veridical. Since this is impossible, he concludes that veridical hallucination must be impossible. I defend against Soteriou's argument by giving an externalist analysis of the case he wields against the possibility of veridical hallucination. In particular, I argue that the perceptual facts make a proto-semantic difference to the contents of visual experiences. That is, I argue that the perceptual facts are among the constituting determinants of the contents of visual experiences. Whether you are the victim of a hallucination or not (and whether you are seeing one thing or another) can make a difference to what must be so if your experience is to be veridical.

The second half of Chapter 3 is devoted to assessing the strength of John Searle's anti-externalist descriptivist account of the contents of visual experiences. According to Searle, visual experiences represent not only distal states of affairs – e.g. that there is a blue thing at location L – but also the obtaining of a causal dependence between those distal states of affairs and the appearance that they obtain. (Searle 1983: 48) Thus, he thinks that perception-grounding relations are represented in the contents of visual experiences. If Searle's account is correct, then the Gricean arguments must be rejected. However, as I argue, there is no reason to accept Searle's account that is sufficiently powerful to outweigh the appearance that veridical hallucination is a possibility. My treatment of Searle has two main parts. In the first, I examine the connections between content and phenomenology. It turns out they are not sufficient to defeat Searle's account, but neither do they lend it strength. Thus, they do not provide sufficient warrant for Searle's account to reject the possibility of veridical hallucination. Additionally, I provide an externalist argument that uses a connection between perceptual facts and agency to reach the conclusion that the perceptual facts are difference makers not only for

the contents of visual experiences, but also for the contents of intentions. The case at the heart of this argument provides another opportunity for the critique of Searle. It turns out that Searle cannot deploy his anti-externalist descriptivism in this case. It looks like he will have to be stuck with the kind of conclusion the rejection of which motivates his anti-externalist descriptivism. Thus, again, we have no reason to prefer Searle's account of the contents of visual experiences over the possibility of veridical hallucination.

2.0: THE STANDARD VIEW, EXTERNALISM, AND VISUAL CONTENTS

Imagine a case of seeing a thing. According to the standard view (herein *TSV*) seeing a thing involves several elements: the particular thing seen, the experience we enjoy in seeing the thing, and a dependence relation between the thing and the experience of the thing. What differentiates TSV from my account lies in the relation it postulates between visual contents and the perceptual facts. I affirm, while TSV denies, that the perceptual facts make a difference to the contents of visual experiences.²

When we see a particular, our experience is experience of that particular, but it is possible that experience with the very same content/phenomenal character be of another thing, or nothing at all. Imagine being enveloped in a self-luminous pink haze. With your eyes open you can see the haze. Your experience bears on the particular pink haze in which you are enveloped. Now imagine being enveloped in a second, but qualitatively identical, haze. This second haze would produce an experience qualitatively identical to the experience you enjoyed in seeing the first haze. Moreover, it is possible that you

² See (Davies 1997) and (McGinn 1989) for the *ur*-statements of TSV. It is probably worth stating that the basic picture – according to which our experiences might possess a particular content and phenomenal character no matter what the perceptual facts – is quite ancient. Descartes, for example, seems to affirm such a view in the Second Meditation.

could, perhaps being overcome by intoxicating fumes, hallucinate a self-luminous pink haze. In all three cases, you enjoy experience with the same phenomenal character, but in each case there are differences regarding the of-ness of the experience. In the first case, you see the first haze. In the second case, you see the second haze. In the third case, you see nothing at all. Your experience does not bear on any existing thing.

Thus, on TSV, whether or not our visual experience bears on a particular object is not determined by its phenomenal character. Given, this line of thought continues, that the content of visual experience is a phenomenal concept, it follows that, when we perceive a thing by entertaining an experience, it is possible that an experience with the same contents could be an experience of another thing, or nothing at all. Thus it is proposed that visual experiences have what might be called *existentially quantified contents*. Such contents are perspicuously represented using devices of existential quantification to indicate those objects they are about.³

There is much that is correct about TSV. In particular, TSV is right about the basic constituents of the seeing relation: seeing a particular is, as I will argue, constituted by its being true that a visual experience bears the right dependence relation to a particular. However, as I argue in Chapter 3, it is false that the contents of a perceptual visual experience must be identical to the contents of its indistinguishable cousins – those experiences which would be just like it, but constitute perceptions of other particulars or nothing at all.

According to TSV, when we say that an experience is, for example, an experience of Sally, its being of her can not be supposed to be entailed by its content. The of-ness of

³ More discussion of the (somewhat obscure) claim that visual experiences have existentially quantified contents is found in Chapter 3.

the experience is not part of its content. The bearing of a perceptual experience on its objects is not part of its content. This might be thought to be an immediate problem for TSV, but it is not. The bearing of a representation on a particular need not be reflected in its content. Thus, for example, you might announce to your kindergarten class that *some dirty little boy has peed on the floor in the coat room* and thereby speak about Sam, but that Sam is who you are speaking about is not determined by the content of your utterance. After all, it is consistent with what you say that Kenny is the perpetrator of the crime. That you have spoken about Sam is a matter of implicature and, to the extent your audience picks up on your sly accusation of Sam, it is because they possess information you have not asserted.

According to TSV, when we say that an experience is, for example, an experience of Sally, we cannot correctly mean that its being *of* her is entailed by its content. Probably, the advocate of TSV would claim that what earns our right to describe an experience as of something is its being part of an episode of perception. Thus the difference between me and the advocate of TSV can seem quit slim here – it is the difference between claiming that the of-ness of a perceptual experience is reflected in its content (and the fact that it is part of the content of experience is constitutively explained by the perceptual facts) and thinking that it is not (but that the aptness of describing a perceptual experience as of something is explained by the perceptual facts).

Indeed, but for the case for externalism that I describe in Chapter 2, I'd be an advocate of TSV. Thus, while I agree that TSV has gone wrong in affirming that the perceptual facts make no difference to the contents of visual experiences, I also think that much of the criticism of TSV has been misplaced. There are two main lines of criticism I

have in mind here. First, there are worries about whether TSV captures the particular-ish phenomenal character of visual experience. Second, John Campbell argues in Chapter 6 of his book, *Reference and Consciousness*, that TSV cannot be right. (Campbell 2002) Campbell claims that if TSV is correct, then it is mysterious how my seeing a thing should put me in a position to understand demonstrative reference to it.

Let me turn to the first worry. In a recent essay, Chalmers gives the following argument (though whether he endorses it with full certainty is not clear) that TSV must be false:

“[A] merely existential characterization of phenomenal content does not fully respect the directness of an experience of an object...experience does not merely present that there is an object at a certain location with a certain color: it presents that *that* object is at a certain location with a certain color.” (Chalmers 2006: 109)

Peacocke gives a similar argument in a footnote in *Sense and Content*:

[If] we are to be strict, the attribution of a common existential content to [visual] experience is too crude. There is a sense in which, as one incompetently says, a hallucination presents a *particular* non-existent object, and so has more than a general existential content. (Peacocke 1983: 9)

Both of these passages try to identify a problem with the thesis that experiences have merely existential content. Both are pointing to facts about the phenomenology of visual experience – their particular-ish phenomenal character – and taking this as reason to reject TSV in favor of other views. In Chalmers’ case, he proposes to accept Burge’s account, according to which the content of experience “contains a demonstrative element.”⁴ (Burge 1991: 202)

⁴ Now I am not sure whether the suggestion that content of experience is demonstrative is helpful. The problem is that sentences containing demonstratives need not differ in which proposition their assertive utterances express from sentences that do not contain demonstratives. For example, on the standard semantics for demonstratives and names, the sentences ‘Fred is happy’ and ‘He is happy’ (where ‘he’ is used as a demonstrative pronoun) may well possess the same content. Sadly, here is not the place to pursue these questions. As we’ll see in the next paragraph, we can get an ever-so-slight traction on the issue – though this traction gives out almost immediately.

The supposition here that phenomenology determines content is problematic – as we’ll see in Chapter 3. But I’d like to focus on a different worry. Susanna Siegel has recently made a good statement of it:

One way to question this would be to say that Chalmers [and company] are all wrong about the phenomenology – that it really is respected by a merely existential characterization of phenomenal content. Another way to question is to ask what exactly the phenomenal difference is supposed to be between the way Chalmers [and company] take the phenomenology to be, and the way they think their opponents take it to be. I find this pretty obscure. Both demonstratives and certain existentially quantified formula (e.g., the Russellian expansion of ‘the table over there’) are devices for representing particular objects. Of course there are modal differences, but those seem irrelevant when considering phenomenology. (Siegel MS)

Siegel is right about this: proposals about the contents of experience – especially proposals designed to capture the bearing of experience on particulars – turn out, often enough, to be modal claims. Thus for example, on standard semantics for demonstratives, if we join Burge in thinking that the bearing of experiences on their objects is demonstrative, presumably we mean to indicate two things: (i) that their bearing on them is direct and (ii) that their bearing on them is rigid.^{5,6}

One might think it would be a simple matter to resolve this question. After all, what is at issue here is the phenomenology of experience. Presumably, therefore, in order to determine the modal profile of the particular-oriented intentionality of

⁵ It is plausible that we might also mean to indicate that, in addition to possessing a content, such experiences possess a character.

⁶ The claims that experiences are rigid and directly referential mean the following: An experience would be rigid with respect to, for example, Fred, if and only if its content is such that its veridicality (at possible worlds in which Fred exists) depends on how Fred is in that world. An experience would be directly referential to, for example, Fred, if and only if, its bearing on Fred were non-descriptive. That an experience’s bearing on Fred is non-descriptive means that it represents Fred and not only by representing the obtaining of conditions which Fred, in fact, satisfies. Thus, we can see the point of the talk of directness. If an experience is not directly referential, then its representing of Fred is mediated by its first representing conditions Fred satisfies. It is sometimes suppose that if a representation is directly referential, then it must be rigid with respect to what it directly references. (Soames 2002: 264) This is a mistake, though arguing so is beyond the scope of this dissertation.

experience we'd simply need to engage in an introspective project. Look at your hand. Now ask of the conscious episode you thereby enjoy whether the way it presents the world to be would be satisfied at different situations. If things turned out this way or that, would your experience be veridical? But this procedure seems hopeless. It doesn't seem like introspection has anything to say about the modal profile of the bearing of our experiences on their objects.

Part of the problem is this: we just do not have much occasion to make fine discriminations between the contents of experiences. Moreover, the aspect of content we're after – its particularity – is presumably widespread. Maybe every visual experience has this feature. As a result, it is not as if we have occasion for discriminating among different experiences regarding their particularity.

The dialectical situation is different in the case of other intentional states. In the case of beliefs and statements of belief, for example, philosophers are trained to make fine distinctions. So we can be expected to do better in our theoretical understanding of belief contents. Not only that, but we have techniques for detecting what might seem to be fine distinctions. We can test proposals about the contents of belief by looking into the web of inferential relations into which beliefs fall. Does this proposal about the contents of some class of belief get the inferential behavior of these beliefs right? Does it forbid the invalid inferences and allow the valid ones? Does it in some sense explain the validity of the valid ones? Another source of data for semantic proposals comes from embedding the sentences that might express belief in modal contexts. So, for example, Kripke argues against descriptivist proposals by showing that if they are right, some

sentences involving names under the scope of the necessity operator count as true, when they ought to be false.

We cannot do this in the case of experience. Part of how we get a grip on the contents of beliefs by getting a grip on how sentences expressing them function. The embedding tests, for example, require that we have access to the syntax of sentences. We reach an understanding of the inferential properties of a content by putting together derivations and contemplating them. Obviously, we cannot embed visual experiences under modal operators and then check to see if we get the right result.

Thus, it seems we cannot directly check the proposal that experiences have demonstrative contents (supposing we understand what this means) rather than merely existential contents against phenomenological data. So the advocate of TSV need not be moved by Chalmers and Peacocke's worry and neither should the critic of TSV rely on it.

Let me turn now to discussing Campbell's criticism of TSV. Campbell's criticism is that if TSV were true, then we could not make sense of our understanding of demonstrative expressions. Imagine the following dialogue: My friend says "That apple is a yummy looking one". I wonder what object she is talking about and she directs my gaze toward the apple in the bowl. Having seen it, I am now in a position to agree or disagree or otherwise engage her in conversation about the apple. My perception of the apple explains the knowledge of reference which explains my understanding of her speech.

Moreover, the same kind of knowledge probably plays a role on the other side of the speaker/hearer divide. Prior to my coming to understand her speech, my friend's perception of the apple equipped her with the cognitive wherewithal to form the

appropriate communicative intentions – to bear the apple in mind in such a way that she could speak with knowledge of reference. Our perception equips us with the same ability, an ability which gets expressed both in speech and understanding of speech. So far so good.

Campbell focuses in particular on the case of demonstratives and argues that conscious experience provides just the sort of window on the world we need in order to engage in knowing dialogue involving demonstrative terms. Thus, perception provides knowledge of reference because it involves conscious experience.⁷ He then uses this connection between experience and knowledge of reference to argue against TSV and in favor of his own view.⁸

According to TSV, when we say that an experience is, for example, an experience of Sally, we cannot correctly mean that its being *of* her is entailed by its content. Thus, that I am perceiving one thing rather than another, or nothing at all, is not given in my conscious experience. I can receive no guarantee, just by consulting the character of my experience that, for example, a particular object exists. I can receive no guarantee that when I wish to direct my visual attention on something that I'll succeed. My attempt to do so can fail because there can be nothing to attend to. My attempt can fail because I can attend to one thing when I wished to attend to another. My attempt could fail because it is pre-empted by hallucination.

⁷ Of course, it provides such knowledge only against a background in which the necessary linguistic and cognitive competencies and motivations are already in place. Thus, when Campbell speaks of a certain kind of mental state as insufficient for grounding an understanding of demonstrative reference, he means that it is insufficient, even given the necessary background. I'll adopt this convention as well in what follows, suppressing explicit reference to the necessary background.

⁸ His view is discussed below in the section on pages 14 and 15.

Campbell argues that if consulting the character of my consciousness won't guarantee these kinds of cognitive success, then we do not know what we're talking about when we use demonstrative expressions. Since knowledge of reference is crucial to understanding our use of demonstratives, our use of demonstratives is mysterious if TSV is right. Or, put another way, since we often do know what we're talking about when we engage in discourse using demonstrative expressions, TSV must be false.

Campbell claims that if your sensory consciousness fails to settle the question of what you are seeing (if anything), then "perception could not provide you with knowledge of which particular thing you are talking about. ...[No] amount of looking to see could provide you with knowledge of which cup is being talked about." (Campbell 2002: 124-25) Therefore, if TSV is correct, our imagined dialogue is a mystery. It cannot explain how we know what we're talking about.

Campbell seems to understand settling the question of what you're seeing (if anything) as working something like this: Attend to one's sensory consciousness and thereby receive a guaranteed answer. The question-settling has to be made possible by something subjectively available. (Campbell 2002: 125) Of course, on TSV, experience is subjectively available, but does not possess a guaranteed bearing on a particular object. Despite my best efforts to perceive, I might fail to see a given individual. Moreover, that I am subject to experience with the same content and phenomenal character as a perceptual experience is consistent with my seeing nothing at all.

Campbell has required that whatever grounds our knowledge of reference must be such as to guarantee an answer to the question of what you're seeing (if anything). Thus, Campbell advocates a view according to which the perceptual facts are constitutive of

experiences. (Campbell 2002: Ch. 6) Experiences consist in the tokening of acquaintance relations with worldly objects, their features, and their relations to the subject of experience. Campbell calls this account of the metaphysics of experience the *Relational View*. On the Relational View, if an experience is a perceptual experience, then it is necessarily perceptual.⁹

However, Campbell's argument against TSV does not succeed. He has nothing convincing to say about why the guarantee – that which constitutes our knowledge of the reference of a demonstrative – must be located in the experience rather than in the way the experience is embedded in the environment. Consequently, the advocate of TSV might agree with Campbell that perception can provide knowledge of reference, but would argue that the guarantee Campbell seeks could be provided by the perceiving relation and not by the metaphysical constitution of the experiences involved in perception.

What reason does Campbell have for thinking that the guarantee he seeks must be located in the constitution of conscious experience, rather than in relations between experiences and the environment? Campbell is not very explicit about this claim, but it seems that he intends his discussion of blindsight to provide some warrant. Campbell argues that a someone with blindsight cannot understand demonstrative reference with respect to objects in his blind spot – at least not as we do.¹⁰ (Campbell 2002: 7-9) What

⁹ Again, there is more discussion of the Relational View (and accounts like it) in the next section.

¹⁰ Someone with blind sight suffers from a dysfunction in his visual system the consequence of which is that, while it is as if he is blind with respect to some portion of his visual field, it is evident that he is nonetheless receiving information about objects and their features in the blind region. People with blindsight can, for example, reach out and configure their hands in such a way as to grasp objects in their blind region. (Milner and Goodale 1995)

provides our knowledge of reference of demonstrative expressions, then, is *conscious* experience of demonstrated or potentially demonstrated objects.

Suppose it is insufficient for understanding demonstrative expressions in the way that we do that one subliminally perceive an object – that one be subject to a non-conscious experience that bears the perception-grounding relation to an object. This is still not enough to establish that the guarantee given by perception must be determined by constitution of the conscious perceptual experience, rather than the relations it bears to objects insofar as it is perceptual. Why is it important for our understanding of demonstrative reference to a thing that we non-subliminally – that is, consciously – see a thing? An advocate of TSV might argue as follows: The reason conscious perception is necessary for our understanding of demonstrative reference is because only conscious perception *makes available to the subject* the guarantee that is constituted by perception-grounding relations. Someone who sees a thing subliminally is like a job seeker whose job offer – his guarantee of employment – has been mailed but not yet received.

So just as the advocate of TSV need not be moved by Chalmers' and Peacocke's worry and neither should the critic of TSV rely on it, she might also reasonably resist Campbell's argument that the Relational View is to be preferred to TSV.

3.0: THE INSURGENT VIEW: DISJUNCTIVISM

My view has some kinship with the main insurgency against TSV and the purpose of this section is to distinguish myself from the main body of the insurgency, disjunctivism.¹¹ Disjunctivists characteristically claim that there is no common mental

¹¹ (Hinton 1967) is the contemporary source for disjunctivism, followed by (Snowdon 1980/81), (Snowdon 1990), and (McDowell 1982). Its current champion is Mike Martin. (Martin 1997) (Martin 2002) (Martin

kind formed by even subjectively indistinguishable perceptions and hallucinations. As

William Child puts it:

[T]he idea of experience, neutral between vision and hallucination, is an abstraction from the more specific and fundamental ideas of vision and hallucination. On this view, then, there is a single sort of *characterization* which can be applied to cases of vision and to cases of hallucination; in both, it looks to S as if something is F. But there is no single type of *state of affairs* common to cases of vision and hallucination. Rather, any case of its looking to S as if something is F will be more fundamentally, and specifically, *either* a case of something's looking F to S [in case perception is occurring], *or else* a case of its merely being for S as if something looked F to her [in case hallucination is occurring]. (Child 1994:144)

Thus, while it can be accurate to characterize episodes of perception and episodes of hallucination in the same terms, our doing so is made accurate not by one underlying mental reality – a visual experience – that is shared between the two but by two fundamentally different mental states. This is a bit obscure, but Child provides a nice analogy:

[C]onsider the two concepts 'photograph of S' and 'drawing of S', and the concept 'likeness of S'. In this case, it is clearly correct to treat the concept of 'likeness of S' disjunctively. Any particular likeness of S is fundamentally, and more specifically, *either* a photograph of S *or* a drawing of S; it is a likeness of S only derivatively – in virtue of being *either* a photograph *or* a drawing. (Child 1994: 145)

Thus, just as it would be a mistake to suppose that the difference between a photograph of S and a drawing of S could be explained by a common thing (a likeness) with different causal histories in each case, it would be a mistake to suppose that the difference between a perceptual experience and a hallucinatory experience could be explained by invoking one thing – a visual experience – with different causal histories.

2004) (Martin 2006) Jonathon Dancy also presents disjunctivism in his article, 'Arguments from Illusion'. (Dancy 1995)

The characteristic claim of Disjunctivism should probably be understood as a consequence of their characteristic metaphysics of perceptual visual experience. Thus, though disjunctivism gets characterized as above, its primary commitment is to a view of the metaphysics of perceptual visual experience much like Campbell's Relational View.

Mike Martin describes the disjunctivist metaphysics as follows:

According to [the disjunctivist metaphysics], the actual objects of perception, the external things such as trees, tables, and rainbows, which one can perceive, and the properties they manifest to one when perceived, partly constitute one's conscious experience, and hence determine the phenomenal character of one's experience. (Martin 1997: 84)

What is the meaning of this talk of constitution? It is supposed to indicate that these objects and their features are essential to the identity of the experiences constituted by them. Thus, the experience I am having now – assuming it is perfectly perceptual – could not exist in worlds in which it is not part of an episode of perception. Were I hallucinating, a subjectively indiscernible, but numerically distinct, experience would take its place. Since purely hallucinatory experiences have no objects, they could not be constituted by objects of perception and their features. Thus, the disjunctivist metaphysics leads almost directly to disjunctivism's characteristic claim. If this metaphysics is correct, then episodes of perception and episodes of hallucination must have rather different constitutions and it is no surprise they do not form a fundamental common kind.

What is the mode of constitution by which perceptual objects and their features constitute perceptual visual experiences? Alston and Langsam propose what is called the *Theory of Appearing*, according to which experience consists, most fundamentally, in the tokening of appearing relations – construed as a relation taking a subject of experience,

an object, and at least one property. (Alston 1999) (Langsam 1997) The relation of appearing is the relation that obtains between, for example, myself, a ball, and redness when a ball appears red to me. The theory requires that perceptual experiences fundamentally consist in relations between subjects and worldly objects – things like sandwiches, people, cars, and rainbows. That perceptual experience *fundamentally* consists in such relations means that the appearance relation is not to be constitutively explained by anything. It is a primitive relation and possesses no constitutive explanation. Thus, according to the Theory of Appearing, experience is not to be understood in terms of relations to contents. A burrito's appearing tubular to me is not to be understood (even partly) as a matter of being in some state with tubularity-involving propositional content – even if it should turn out that experiences have propositional contents.

It is not my intention to present an assessment of disjunctivism in this dissertation. There are already fairly convincing critiques of the view.¹² My goal is to distinguish myself from the disjunctivist. Like the disjunctivist, I think that there are differences (philosophically interesting differences, deep differences!) between perceptual experiences and their counterparts. However, I deny both the disjunctivist metaphysics and the characteristic claim of disjunctivism. Until I see a good reason to think it, I'll suppose that perceptual experiences and their counterparts are of the same fundamental kind. When Child claims that it would be a mistake to suppose that the distinction between a perceptual experience and a hallucinatory experience could be explained by invoking one thing – a visual experience – with different causal histories in two different

¹² See (Comesaña 2005), (Wright 2002), (Johnston 2006), and (Hawthorne and Kovakovich 2005) for some examples.

circumstances, I think the proof is in the pudding and we must see how such an explanation might proceed. And, of course, this is the task partially accomplished in Chapter 2. Thus, while I think the perceptual facts make a difference to the contents of visual experiences, I do not think that this supports the view that perceptual experiences do not form a common kind with hallucinatory experiences.

CHAPTER 2

PERCEPTUAL DEPENDENCE: FINDING A PLACE FOR COUNTERFACTS

The business of this chapter is to present and advocate for my account of what constitutes a visual experience as a perceptual experience – what constitutes an episode as an episode of seeing. I will begin by presenting my view. I'll do this by critiquing David Lewis' counterfactual theory – the theory of which my account is an extension and revision. My account is a descendant of Lewis', but it is not the only one. Accordingly, I will present and critique Bruce LeCatt's theory, which, like mine and Lewis', deploys the apparatus of counterfactual dependence to give an account of perception. Next, I will present and critique arguments – arising from the work of Brian McLaughlin – against the thesis that perception should receive a counterfactual analysis of any sort. After that, I will address one of Michael Tye's counterexamples to a counterfactual analysis by invoking analogies to photographic representations. Next I address the possible eliminability of reference to causation in the analysis of perception. Finally, I wrap up with a discussion of objections to a process-relative counterfactual account that arise from my use of the notion of a process.

I favor an analysis that follows David Lewis in including counterfactual dependence in a refinement of Grice's causal dependence theory. In particular, on my view, it is necessary that a perceptual experience be the outcome of a process that supports a suitable pattern of counterfactuals. This differs from Lewis insofar as it

requires that the relevant counterfactuals hold fixed some actual process that links an experience and its objects. My account also differs from Lewis' insofar as it is an account of *transitive* seeing. That is, I am interested in the question of what makes it the case that I see one thing, rather than another. What makes it true, for example, that in looking out just now, I see this cup rather than the one I could see if I were to glance to the right? Lewis' theory is designed to account for *intransitive* seeing – what makes it true that I see anything, rather than nothing at all. What makes it true that I see? Full stop. Thus my account is a refinement of Lewis' but it is also an extension of it.

I should add that Lewis' division between accounts of transitive seeing and accounts of intransitive seeing is problematic. The reason is this: every instance of intransitive seeing is also an instance of transitive seeing. I cannot see without seeing something. Thus, one would expect that the way to give an account of intransitive seeing would be to give an account of transitive seeing. Whatever makes it true that I see something is also what makes it true that I see.

My refinement of the counterfactual account is attractive because it avoids the many counterexamples that have been marshaled against the counterfactual accounts of perception, while preserving what is attractive about counterfactual accounts.

PART 1 COUNTERFACTUAL DEPENDENCE AND SEEING

1.1: COUNTERFACTUAL DEPENDENCE AND SEEING

My proposal arises from the dialectic begun with Grice, who argues that it is necessary that a perceptual experience (an experience that is part of an episode of seeing)

causally depend on what is seen.¹³ (Grice 1961) His two main arguments for this view are the Argument from Veridical Hallucination and the Argument from Competing Objects.¹⁴ Moreover, Grice argues that it is sufficient for seeing a thing that one's experience causally depend on it in the right kind of way. (Grice 1961) Thus, on Grice's view, it is necessary and sufficient for seeing a thing that one's experience be caused by the thing in the right kind of way.

Grice's motivation for specifying that perceived object must cause experiences in the right kind of way stems from what are now entirely familiar problems. Mere causal dependence won't do. There are two problems here: the problem of causal deviance and the selection problem.

Cases of deviant causation tell us that mere causal connection is insufficient for seeing a thing. Consider the following case:

Raven Case: A neurosurgeon is manipulating your brain while you are conscious. The point is to electrically stimulate your visual cortex and find out what sorts of experience you have. Suddenly, a raven flies into the room and attacks the surgeon. In his haste to escape the bird he accidentally stimulates your visual cortex in just the right way for you to have experience which is just like it would be if you were actually seeing the raven.

In the Raven Case you do not see the raven. Thus, even though your experience is caused by the raven, you do not see the raven in being subject to the experience. A mere causal dependence between a visual experience and an object is therefore insufficient for

¹³ Grice himself was responding to H.H. Price's theory. However, Price's work belongs in the prehistory of this dialectic, since Price was concerned with establishing that perceptual consciousness is a form of inference, in particular a form of causal inference. (Price 1932: 395) Thus, on Price's account, perception is essentially a way of coming to believe something. Like Lewis and Grice, I understand perception to be a way of coming to have an experiential state. Experiential states differ from beliefs in many ways. For example, we hold people responsible for their beliefs, but not, given the direction of their gaze, for the character of their experience.

¹⁴ I discuss these extensively in Chapter 3.

seeing a thing. Thus, in order to achieve sufficiency for a causal condition on seeing, Grice adds that the experience must causally depend in the right way.

Let me explain the selection problem. Suppose that Fred is Zenobia's grandson. Suppose that I see Fred, but not Zenobia, who is far away in California. But, both Fred and Zenobia are part of the causal chain leading to my experience. Since I see one and not the other, it follows that mere causal dependence is not enough for perception.

Mere causal dependence is too permissive: it generates the selection problem and the problem of causal deviance. Thus, if a causal account of perception is going to work, it had better place restrictions the kinds of causal relationships that might partly constitute an episode as an episode of perception. Thus, Grice claims that the causal relationship between an experience and perceived object must be the right kind of causal relationship. What, on Grice's account, is the right kind of causal dependence? Grice supposes that the best way for the theorist to proceed is to indicate, by means of examples, what kind of causal dependence is required. Thus, it is sufficient for X to perceive an object that:

it should be causally involved in the generation of some sense-impression by X in the kind of way in which, for example, when I look at my hand in good light, my hand is causally responsible for its looking to me as if there were a hand before me, or in which ... (and so on), whatever that kind of way may be; and to be enlightened on that question one must have recourse to the specialist. (Grice 1961)

On Grice's account, there is a kind of gap in our concept of perception to be filled in by the scientist. The situation here is a bit like the situation of the folk with respect to material-stuff concepts. Water is the kind of stuff that is made of whatever the scientific authorities tell us it is made of. Likewise, Grice is urging, perception is a causal process the essential character of which is to be explicitly determined by the scientist.

Grice supposes that whatever fills this gap will also solve the problems of selection and causal deviance. Of course, even if we accept that our conception of perception has a gap to be filled in, one might hope to use philosophical methods to say more about how the gap might be filled in. We can, therefore, see Lewis as contributing to this project. In his article, 'Prosthetic Vision and Veridical Hallucination', David Lewis proposes the following dependence condition: "If the scene before the eyes causes matching visual experience as part of a suitable pattern of counterfactual dependence, then the subject sees; if the scene before the eyes causes matching visual experience without a suitable pattern of counterfactual dependence, then the subject does not see." (Lewis 1986a: 281)

The kind of counterfactual dependence relation Lewis imagines here is just the kind that he thinks reductively grounds causal dependence. (1986a: 166-67) (Lewis 1986a: 284) If some scene S and some experience E bear the right kind of counterfactual dependence, then two things are true:

- i) if S were to occur then E would occur;
- and
- ii) if S were not to occur then E would not occur.

Additionally, the occurrence of E must be part of a *suitable* pattern of dependence. There are, on Lewis' account, two aspects to this suitability:

First, if things were suitably different with respect to how they look, then the different scene would cause suitably different visual experience. (Lewis 1986a: 286). The suitability in question seems to be a matter of which differences (properties) count for the pattern of counterfactual dependence. It won't do to claim that no one sees

anything because, for example, were things different in invisible respects no experience would change. Moreover, it won't do to claim that no one sees anything because suitable changes in the environment wouldn't cause experiences which vary in respect to, for example, their time of occurrence. The idea is that suitable changes in the environment need to be accompanied by changes in how experiences represent the environment. Thus, for example, if a white wall were to differ in respect to its color (e.g. by being vividly orange), then that difference would be accompanied by an experience which represents the wall as being orange.¹⁵

Second, if some scene S and experience E are related as part of a suitable pattern of dependence, then that pattern needs to pair a large enough set of suitably differing scenes with a large enough set of differing experiences. That is, there must be *plenty* of alternative scenes which differ in *plenty* of visible respects such that, for each one, if it were to occur, an experience which matches it would also occur, and if it were not to occur that experience would also not occur.¹⁶ (Lewis 1986a: 282-83)

Putting things together: To see is for your experience to track possible changes in your environment. What determines whether or not you see in a situation is things being *just so* in modal space. The key to Lewis's position is this: His requirement that an experience be linked to its objects by a suitable pattern of counterfactual dependence is a way of cashing out the thought that perception must be an expression of a capacity to discriminate by sight. (Lewis 1986a: 286) Lewis thinks of discrimination as a function from possible environments to experiences that truly represent those environments.

¹⁵ Lewis thinks that it is an empirical question which features matter here. (Lewis 1986a: 282)

¹⁶ Lewis denies that 'plenty' needs to be strictly defined. It may be a vague matter whether a person counts as seeing in a given situation. The uncertainty about whether a person counts as seeing in a given situation depends on the way in which the domain of required matching differs in its size in different contexts. (Lewis 1986a: 283)

Thus, one exercises a capacity for discrimination just in case some appropriate set of counterfactual relations obtains.

How would adverting to counterfactual dependence help us solve the problems of selection and deviant causation? Cases of deviant causation tell us that mere causal connection is insufficient for seeing a thing, but they also prompt the hypothesis that counterfactual dependence matters to seeing it. This is because a theory of counterfactual dependence has the resources to explain why, though your experience causally depends on it, you do not see the raven. The process linking you to the raven is such that, even if it were operational, were the raven different in relevant ways your experience wouldn't track those differences. The process, after all, has the panic of the neurosurgeon as part of it. The fact that the panicked doctor stimulated your brain just as he did is chancy. It just isn't true, for example, that if the raven were blue, the doctor's panicked actions would produce experience as of a blue raven. Consequently, the right pattern of counterfactual dependence does not obtain. Counterfactual dependence provides an elegant explanation of the fact that you do not see in the raven case.

Let's turn now to the selection problem. I want to know: given that Fred and Zenobia are both part of the causal history of my experience, why do I see one and not the other? How shall we design a theory to select from among *all* the particulars playing a role in the causal history of an experience *only* the perceived ones? Again counterfactual dependence provides an elegant solution: you see individuals at stages of the causal chain if and only if your experience appropriately tracks possible changes in them.

1.2: THE CENSOR CASES: CONTRA LEWIS

Despite the appeal of a counterfactual theory of perception, there are counterexamples that pose a serious challenge to Lewis' theory – i.e. censor cases. These are cases where an inactive censor waits off stage, as it were, never actually acting, but disposed to act in such a way that Lewis's preferred patterns of counterfactual dependence are disrupted. Insofar as these censor cases are cases of perception, Lewis's counterfactual dependence theory is threatened. If one sees in a censor case, then the counterfactual dependence of the sort Lewis favors cannot be a necessary condition on being a perceptual experience. Consider, therefore, the following case:

The Censor: Imagine a supernaturally powerful being – call him the *censor* – who wishes that you be subject at a particular time to experience with a particular intrinsic phenomenal character or predicative content. Maybe, for example, he wants to be as if you were looking at a white wall. As it happens, the censor remains inactive. His desires are realized by the actual course of events. At the time during which he wants for it to seem as if you were looking at a white wall, you are looking at a white wall and everything functions as it ought to.

But, were visible events different, the censor would step in and produce the experience he wants you to have. If you were standing in front of a blue wall, or a grey hippo, or whatever, he would intervene and produce in you the experience you actually have. He would pre-empt the nascent causal chain stemming from those events and cause in you the experience he wishes you. Consequently, the censor makes it false that were things different with the visible environment your experience would be appropriately different.¹⁷

The dominant intuitive response to the story of the censor is that you see, despite the existence of this censor. But, if Lewis is right about the counterfactual dependence condition on perception, you cannot see in the censor cases. This is just because were things different in the relevant ways, your experience wouldn't track those differences.

¹⁷ This kind of case shows up in Bruce LeCatt's article, "Censored Vision". (LeCatt: 1980). Interestingly, there is abundant reason to think that Bruce LeCatt was written by none other than David Lewis' cat. More on LeCatt's theory below.

Thus, either Lewis' theory is mistaken, or the dominant intuitive response is confused. Which is it?

1.3: DO YOU SEE IN CENSOR CASES? LEWIS' ASSAULT ON COMMON SENSE

Lewis has provided some reason to doubt that seeing is going on in the censor cases. According to Lewis, our disposition to judge that perception has occurred is based on a misunderstanding of how to compare the censor cases to exactly similar cases in which the censor is absent. Since the censor remains inactive, it is tempting to think that the experience in these cases is perceptual, as it would be if there were no censor at all. Things actually proceed as they ought to. This temptation rests on the presumption that the censor's inactivity is extraneous to the explanation of the connection between the experience and the scene that causes it.

Lewis argues that this presumption is a mistake. (Lewis 1986a: 286) The censor is, the argument goes, an essential part of the situation. Lewis invokes the example of a logic circuit in order to establish that it is important that we take into account inactive, but potential causes when engaging in causal reasoning: "Think of a circuit built up from exclusive-or-gates: every output signal from such a gate is caused partly by the absences of a second input signal." (Lewis 1986a: 286) In such a case, it seems pretty clear that the potential but non-actualized signal plays a part in the causal explanation of the circuit's output. Lewis uses this to bolster the claim that the censor cases are quite different from ordinary cases. On Lewis' analysis, the presence of the censor is an essential feature of the censor cases and it is, of course, a feature missing in ordinary

cases. So, Lewis urges, they cannot be very similar after all and we should not generalize from censor-free cases to censor cases.

But the example of the logic circuit only suggests that in *some* contexts non-active (but potential causes) are relevant to the explanation of occurrences.¹⁸ Lewis has not established that the presence of a non-active censor is sufficient to make it illegitimate to apply the concept of perception in the censor cases. The causal explanation of logic circuits may require taking into account non-active (but potential causes), but this is probably not true in the censor cases.

It is part of the proper functioning of the circuits Lewis describes that their output depend not only on what signals pass through them, but also on what signals do not pass through them but could have. We have designed the circuits like this and so it makes sense that the explanation of their activity should take into account non-active but potential signals. Thus we have justification for thinking that non-actual (but potential) causes are part of the causal explanation of the activity of the logic circuit. However, it is not obviously part of an explanation of the proper functioning of the visual system that a censor monitors and influences its activities. There is a principled reason to treat cases like the censor in a way that differs from how we treat cases like logic circuits. The reason Lewis gives for rejecting the comparison between censor cases and their censor-free counterparts may be side-stepped. Lewis has not succeeded in arguing that intuition leads us astray here. Thus, these cases stand as un-defused challenges to Lewis' account of the seeing relation.

¹⁸ It is not even clear to me that it is a case where a non-active but potential cause is an essential feature of the situation.

1.4: DO YOU SEE IN SENSOR CASES? AN ANALOGY FROM ACTION THEORY

Lewis thinks that perception is essentially a process by which we discriminate among states of the environment. This is what explains his use of counterfactual dependence. But there is a deeper reason counterfactual dependence is implicated in perception. Perception requires a kind of control. The reason why you do not see the raven in the raven case is that it does not control your experience. In order for you to perceive a thing your experience must be controlled by the thing you perceive. Accordingly, an account of perception invoking causation needs to be supplemented specifically with reference to the control relation.

Consider two objects (Sally and Charlie): Sally controls Charlie if and only if there is some pattern of counterfactual dependence such that states of Charlie track states of Sally. That states of Sally determine states of Charlie constitutes her control over him. If Sally were this way, then Charlie would be this way. If Sally were that way, then Charlie would be that way, etc. Control requires that one set of states be counterfactually depend on another.

Turn now to the kind of control specifically implicated in perception. The kind of control required in order to perceive is the kind of control that makes the tokening of an experience a form of discrimination. Think again about the raven. What was lacking in the case of the raven was a tracking relation between states of the raven and states of your experience, such that states of the raven would be correctly represented by your experience. The raven was not in control of your experience, and, in particular, that lack

of control entailed that your experience was not the expression of a capacity to discriminate among states of the raven.

It should not be surprising that perception requires control. Consider the roles perception plays in our lives. Among other things, it provides a basis for our cognitive and practical engagement with the world. I get in the business of knowing the facts because I perceive things. I successfully negotiate a complex world because I perceive its bits. If control were a central feature of perception, then we could explain how perception plays the role it does. It contributes to our thriving by putting at least some of our acting and thinking under the control of relevant bits of the world.¹⁹ So, by inference to the best explanation, perception requires control.

Given that perception requires control, we can shore up our intuitions about the censor case by drawing analogies between perception and action. Perception requires control, but so does morally responsible action. When I perceive a thing, my experience must be under the control of the seen thing. Likewise: when we act in such a way as to be a candidate for praise or blame, our action must be under our control. A loss of control is among the classic defeaters for claims of responsibility. Presumably, control is implicated in morally responsible action because control is a requirement of intentional action.

One might think, given the connection between control and counterfactual dependence, that if you are morally responsible for an action, there must be a counterfactual dependence relation between your decisions and actions. The accounts of

¹⁹ Of course, perception serves goods other than cognitive and practical ones. I want more from seeing than that I be informed. We often take pleasure in simply seeing things. Accordingly, perception plays a crucial role in aesthetic and, for another example, erotic pleasure. But even here we can make sense out of control's central importance to perception.

perception and morally responsible action will be structurally similar.²⁰ Indeed, just as counterfactual accounts of perception suffer from cases like the censor, a counterfactual account of action would suffer from a similar case. Consider, then, the case of the Neuroscientist's Daughter:

Neuroscientist's Daughter: Chuck is dating Brunhilda (the daughter of Ivan, the neuroscientist). Their relationship is very serious and Chuck is thinking of marrying Brunhilda. Chuck is a fine fellow – handsome, smart, morally upright, and athletic. In short: he is the best mate for Brunhilda. Ivan, the neuroscientist, wants nothing but the best for his daughter and so is pleased to discover that Chuck might marry her. However, Ivan also knows that Chuck is wavering a bit on the question. Brunhilda, wonderful though she is, cusses like a sailor. Ivan knows of Chuck's wavering because he is secretly reading Chuck's diary. Now, Ivan is prepared to do whatever it takes to maneuver Chuck into marrying Brunhilda, but being lazy does not want to take action unless he has to – that is, unless Chuck decides not to marry Brunhilda. As it happens Chuck resolves his doubts, proposes to Brunhilda, she accepts, and the two marry. They live happily ever after.

If Chuck had decided otherwise, then Ivan would have stepped in. Using his vast knowledge of the human brain, Ivan would have made it so that Chuck marries Brunhilda. He would have, maybe, erased the effects of Chuck's decision and put Chuck's brain into the state it would have been in had Chuck decided to marry Brunhilda. So there is *some* sense in which Chuck could not have done otherwise and, yet, intuitively Chuck is responsible for his part of the marriage.

Harry Frankfurt uses cases like this to argue against the Principle of Alternative Possibilities, according to which an agent is morally responsible for having performed an action only if he could have done otherwise.²¹ (Frankfurt 1969) According to Frankfurt,

²⁰ Indeed, the similarities run deep. Just as Gricean considerations motivate a causal theory of perception, Davidson deploys similar considerations in favor of a causal theory of intentional action. (Davidson 1980) Just as problems of deviance and selection arise for perception, the same kinds of problems arise for causal theories of action.

²¹ In the right mood, the Principle of Alternative Possibilities can seem very puzzling. It is a *very* weak requirement and moreover: why should we suppose it is true? What explains the initial plausibility of supposing that being morally responsible for an action requires that one could have done otherwise? It might be that the Principle of Alternative Possibilities follows from the stronger requirement that actions for which you are responsible track one's decisions. Thus, when I act responsibly, the nearby possible worlds at which I decide differently are also worlds where my action reflects the decision I have made. If I have acted responsibly, then in some nearby world, I decide to X rather than Y, and so I do X. Thus, in order to be in control of one's actions, it must be possible to do otherwise. There has to be some world in

since Chuck is responsible for having married Brunhilda, the Principle of Alternative Possibilities must be rejected. That is, Frankfurt takes cases like this to be evidence that you can be responsible for an action, even if you couldn't have done otherwise.

Is Chuck morally responsible for wedding Brunhilda? Remember: the neuroscientist never lifts a finger to ensure that Chuck marries his daughter. Chuck carefully and clearly deliberated about what to do. It seems that he is responsible for wedding Brunhilda: his friends can praise him for making the right choice (or blame him as the case may be). He and his wife can appropriately say to each other "How clever we have been to marry each other!" Moreover, Chuck seems genuinely bound by the terms of the marriage contract, as he would not be if he were not responsible for entering into it.

As anyone who has followed the literature on the Frankfurt cases should know, the answers here are pretty contentious.²² It is not entirely clear that how we should – all things considered – decide the matter of Chuck's moral responsibility. It is not clear how the Frankfurt cases function in the context of a debate about the compatibility of moral responsibility and causal determinism. However, for my purposes I need to settle neither the question of Chuck's moral responsibility (plausible as I find the claim that he is responsible) nor the question of the compatibility of moral responsibility and causal determinism.

What matters to me is whether Chuck's assenting to the marriage was an intentional action of his. It seems clear that it is. Moreover, I find it plausible that Chuck was in control of his actions. It was his considerations that moved him to act. They were the causes of his actions and determined their character. To the extent that we judge

which I do otherwise. Thus, we get the Principle of Alternative Possibilities from a counterfactual account of morally responsible action.

²² For an overview of the debate see (Fischer 1998) and (Widerker and McKenna 2002).

Chuck morally responsible, we are probably doing so on the assumption that the facts described do not rule out his actions being intentional actions.

The analogies to the censor cases are clear. We have a non-active, but possible, intervention. We therefore have a failure of counterfactual dependence. But we want to say that nonetheless a certain capacity has been properly expressed. In the one case, despite the failure of counterfactual dependence, Chuck has acted intentionally. In the other case, despite the failure of counterfactual dependence, you have seen the white wall.

1.5: THE REVERSE CENSOR, THE RAY GUN

In addition to censor cases, which pose a challenge to the necessity of a Lewis-style account of perception, there are cases which pose a challenge to the sufficiency of a counterfactual account. I have in mind two kinds of case – the Reverse Censor and the Ray Gun. Both of these stem from Martin Davies' discussion in 'Function in Perception':

The Reverse Censor: My natural or prosthetic eye is not functioning, and it is by more or less spurious (that is, deviant) means that the scene before my eyes causes matching visual experience. But if the scene were any different my visual experience would be correspondingly different. For there is a 'reverse censor' standing by, ready to see to it that I have matching visual experiences whatever the scene might be. So long as the scene before my eyes is the one that is matched by the caused experience the reverse censor does nothing. (Davies 1983: 412-13)

But for the presence of the reverse censor, this would be a clear case of veridical hallucination, however, the reverse censor ensures that Lewis' counterfactual condition is met in this case. Were things relevantly different in the viewer's environment, then his

experience would be suitably different. However, we want to say that no seeing has gone on here. Lewis has, the argument goes, failed to give us a sufficient condition.

The Ray Gun: A prisoner is made temporarily blind by being held in the path of a ray gun which 'jams' the transmission of visual information. It is usual to hallucinate while being blinded in this way, provided that one's eyes are open. The prisoner has a visual experience as of a uniformly red wall five feet in front of the prisoner's eyes, and the presence of the wall is a causal antecedent of the visual experience. The ray gun is a particularly delicate one, and if the scene before the prisoner's eyes were any different the ray gun would not work. In that case the prisoner would be able to see normally. (Davies 1983: 413)

In the case of the Ray Gun, the scene in front of the eyes causes matching experience as part of an appropriate – Lewis-style – pattern of counterfactual dependence. The fragility of the ray gun ensures this. But we do not want to say that perception is occurring. If Lewis is right, then the prisoner sees, yet (we judge) seeing is not happening. Lewis has, the argument goes, failed to give us a sufficient condition.

Again, either seeing is not happening, intuition is getting things right, and we must reject Lewis' view or we must accept Lewis' view and deny that intuition is giving the right verdict here.

1.6: MY VIEW

We have three counterexamples against Lewis' view. Perhaps Lewis' mistake was to give counterfactual dependence pride of place in his account of seeing. However, the wholesale rejection of counterfactual accounts would be hasty at this point. There are ways around the counterexamples.

The censor, reverse censor, and ray gun counterexamples all involve an interesting feature: in the nearby possible worlds, changes to the scene in front of you

change the processes that eventuate in your experience. There is some process operational in the actual situation, but were the scene different, a different process would cause an experience.

In the censor case the censor remains inactive, allowing the usual process to generate veridical experience. However, his intentions are such that were the scene different – were the scene such as to produce in you a different experience – he would intervene. The tracking relation between your experience and the world is foiled, because, *between* the actual situation and the relevant possible situations, there is a change in process.

In the case of reverse censor the same thing is true. Here the pattern of counterfactual dependence obtains because, were the scene different, the process actually operational would be replaced by a different process. The reverse censor would step in to produce the desired experience – in this case, a veridical experience. Similarly, in the case of the ray gun, the fragile ray gun would be broken by changes in the scene and so no longer interrupt the causal processes that would produce in you matching experiences. In the ray gun and reverse censor cases Lewis' tracking relation obtains, but only in virtue of the fact that between the actual situation and the relevant alternatives there is a change in process.

Accordingly, if counterfactual dependence requires that the processes causing your experience remain stable, we no longer have counterexamples. Thus, in determining whether perception has occurred, we should check the subset of the nearby possible worlds in which the process actually eventuating in your experience is operating

to see whether a suitable pattern of changes in the scene in front of you is matched by changes in your experience. Here is my proposal:

I see some object x by having some visual experience e iff there is some process P linking e to x such that for each member, F , of some suitable set of features x could possess (i) were x F and were some visual experience, e_2 , to be the output of P , then e_2 would represent something as F (and, in so doing, increase the number of things that are represented as F by one); and (ii) were x not- F and were some experience, e_3 , the output of P , then e_3 would not represent x as F .

What I've done here is hold fixed the process that connects a perceptual experience and its objects in the counterfactual situations relevant to determining whether seeing has occurred. If this maneuver works, then it will allow us to retain the good things about a counterfactual account, while avoiding the counterexamples.²³

²³ In an earlier draft of this chapter, I gave the following analysis (call it the *First Pass Analysis*) at this point:

I see some object x by having some visual experience e iff there is some process P linking e to x such that for each member, F , of some suitable set of features x could possess (i) were x F and were some visual experience, e_2 , to be the output of P , then e_2 would represent x as F ; and (ii) were x not- F and were some experience, e_3 , the output of P , then e_3 would not represent x as F .

By itself, this analysis is adequate – though, as we'll see later, it could use some further description of the process P . However, when mated to the externalist thesis from Chapter 3 that the contents of experience depend on the perceptual facts, one might worry that clause i is problematic. In particular, the worry is that if an experience's having some object, x , as a representational target depends on the perceptual facts, then clause i covertly presupposes the obtaining of the seeing relation. According to clause i, it is a partial constituent of its being true that one sees x that, in some counterfactual situations, one's experience represents x as bearing a certain property. But, if the externalist thesis is right, then that one's experience represents a thing depends on the perceptual facts. Thus, the worry is that the First Pass Analysis is a viciously circular analysis.

The solution to this worry is to describe the representational content of e_2 in a way that does not presuppose that it possesses a particular object as its representational target. The way to do this is for the analysis to pick out x by description. (Notice that the analysis targets x by description; not necessarily the experience.) We might fail to know whether an experience is a perceptual experience and be consequently unsure which existing object – if any – is its representational target. Nonetheless, we can get some handle on its contents. We could know, for example, that it represents something as being blue. We could know that it represents the very same thing as flat and circular. There is, therefore, a description of its content that abstracts from its bearing on particular objects. We'll need to exploit this kind of description in revising the First Pass Analysis.

It won't do to revise the First Pass Analysis as follows: I see some object x by having some visual experience e iff there is some process P linking e to x such that for each member, F , of some suitable set of features x could possess (i) were x F and were some visual experience, e_2 , to be the output of P , then e_2 would represent *something* as F ; and (ii) were x not- F and were some experience, e_3 , the output of P , then e_3 would not represent x as F . The problem is that, for example, if 'blue' substitutes for 'F', then if there are *other* blue things in your visual field, then your seeing x would not depend on the status of x , but rather

1.7: BRUCE LECATT: STEPWISE DEPENDENCE

Like me, Bruce LeCatt takes the intuitive response that you see in the censor case as given and tries to revise Lewis's theory. (LeCatt 1982) His proposal is as follows: if R is to be a perceptual experience of P then either (i) R counterfactually depends on P or (ii) there is a chain of events linking P to R such that each event counterfactually depends on its ancestor. (LeCatt 1982: 159) In short, the addition allows that there can be a chain of what LeCatt calls 'step-wise counterfactual dependence' connecting P to R.

Thus, LeCatt claims that while R does not counterfactually depend on P in the censor case, R *is* linked to P by a chain of step-wise counterfactual dependence relations. R counterfactually depends on Q and Q counterfactually depends on P – even though R does not counterfactually depend on P. This is what allows R to be a perceptual experience of P.

But how can this be? It is tempting to think that any situation in which two events fail to be counterfactually dependent is also a situation in which they cannot be related by a chain of step-wise dependence. This is because it is tempting to think that counterfactual dependence is transitive. If one event counterfactually depends on another, which in turn depends on a third, then it is tempting to suppose that the first event in the series depends on the last. Suppose that S doesn't happen, since I depends on

on the status of the other blue things. Though I will not argue for it now, this can be expected to make the analysis too permissive.

What we want (in order for seeing to occur) is that, for each member of some suitable set of features, if the thing were to change with respect to it (but not with respect to other members of that set), then that difference would show up in experience in that counterfactual situation. What we want is that differences in x – the thing that might be seen – be difference makers for one's experience. What we want is that were x to change from its actual case, by becoming F (when it wasn't F before), then one's experience would represent a *new* F. Thus I propose the analysis that appears in the main body of the text above.

S, I would not occur and some other intervening event would. But what would happen if I didn't happen? Since E depends on I, E would not have occurred. So it looks to be true if E is step-wise dependent on S then it must also be dependent *simpliciter* on S. The situations that make it false that E depends on S should be the same situations that make it false that E depends on I.

In order to assess the correctness of LeCatt's account we've got to imagine various permutations of the possibilities and determine whether any of them will make the following true at W, the imagined actual world:²⁴

$P \rightarrow Q$
 $\text{not-}P \rightarrow \text{not-}Q$

$Q \rightarrow R$
 $\text{not-}Q \rightarrow \text{not-}R$

The Censor:
 No matter how
 things are with P,
 R would occur.

Each of these places restrictions on how worlds may be ordered in respect to distance from W. If these are all supposed to be true, then there must be some such ordering that simultaneously satisfies the all the restrictions. Of course, if such an ordering is impossible, then LeCatt cannot save Lewis's account by invoking step-wise counterfactual dependence.

It turns out that counterfactual dependence is intransitive. We can imagine cases where two events are step-wise dependent without being dependent *simpliciter*. We might as well turn to what LeCatt is probably thinking of for an example of how this might work.

²⁴ 'P → Q' should be read as 'if it were the case that P occurred then it would be the case that Q occurred'.

Consider the following picture. It represents the ordering of possible worlds which would make it true that P and R are linked by a chain of step-wise dependence, but false that R depends on P:

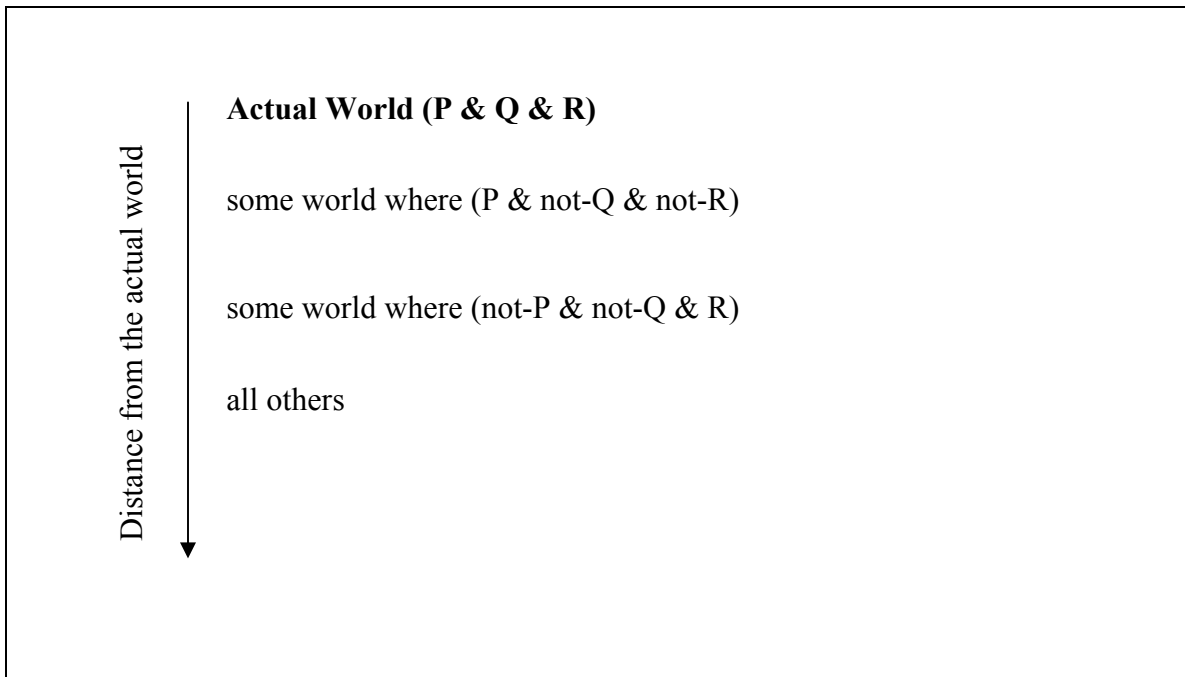


fig. 1 – LeCatt's Censor

Here we see an ordering of worlds that makes it true that $P \rightarrow Q$. First, on the assumption that any world is closer to itself than any other, the actual world ensures that $P \rightarrow Q$ and $Q \rightarrow R$ are true.²⁵ Second, there is some not-P & not-Q world which is closer than any not-P & Q world. This makes it true that $\text{not-P} \rightarrow \text{not-Q}$. Third, there is some not-Q & not-R world which is closer to the actual than any not-Q & R world. This makes

²⁵ Note that it ought to be controversial to suppose that every world is closer to itself than any other world. Lewis thinks that if the actual world is a P & Q world then $P \rightarrow Q$ is vacuously true. This has initial plausibility – especially if we think of closeness as similarity. But it also makes it vacuously false that $\text{not}(P \rightarrow Q)$. It does not *seem* conceptually incoherent to claim that if the relationship between P & Q is indeterministic, then $P \rightarrow Q$ is false. Not only does this seem coherent, it seems true for some substitution instances of P & Q. I will however assume, for this paper, that Lewis is correct about the truth conditions of counterfactual conditionals.

true that $\text{not-Q} \rightarrow \text{not-R}$. Finally, the picture accurately captures the ordering of worlds affected by the activities of the censor. The closest P and not-P worlds are all ones where R occurs. In the censor case, R is not counterfactually dependent on P, but it is R is step-wise dependent on P. And this, LeCatt claims, explains why it is a case of perception.

1.8: THE SUPER CENSOR: *CONTRA* LECATT

LeCatt's theory looks like an improvement over Lewis. However, the type of argument LeCatt employs against Lewis may be used against LeCatt as well. LeCatt uses a censor cases to argue that Lewis's counterfactual dependence theory is excessively strict. But LeCatt has not imagined the most powerful possible censor. Indeed, there is a censor case which may be invoked to suggest that LeCatt's step-wise dependence theory is excessively strict.

Super Censor: God, the omnipotent super censor, wishes you to have experience with a particular character – call it 'R' again. Next, imagine that the scene P occurs (a white wall with blue dots enters your field of vision), next some intervening event Q (some retinal event, say), and finally the experience R occurs. Let me depict the imagined actual course of events like this:

W @ time₁ P... @ time₂ Q... @ time₃ R

Now R is exactly the experience the super censor wants you to have and so He takes no steps to censor your vision. But, again, were things different in respect to how things were in your field of vision, the super censor would step in and ensure that your

Is seeing going on in Super Censor? In such *recherché* cases it is not clear that intuitive judgments can be taken to be authoritative. However, the very same considerations that supported the intuitive judgment that seeing occurred in the case of the Censor. Since the Super Censor remains inactive, we should think that the experience is perceptual, as it would be if there were no censor at all. The Super Censor's inactivity is extraneous to the explanation of the connection between the experience and the scene that causes it.

PART 2 PRINCIPLED OBJECTIONS

2.1: MCLAUGHLIN'S WORRIES

Brian McLaughlin describes counterfactual theories of perception as “deplorable” and urges that they be replaced by an alternative that does not advert to counterfactual dependence. (McLaughlin 1996) The goal of this part is to resist McLaughlin's worries. McLaughlin attributes two views to Lewis: (i) that the exercise of a capacity is constituted by counterfactuals, and (ii) that to perceive is to have an experience as the upshot of the exercise of a capacity for visual discrimination. McLaughlin argues that it is a mistake to think that occurring as part of a suitable pattern of counterfactual dependence is constitutive of exercising an ability to perceive. But he affirms – in agreement with Lewis – that to perceive is to exercise a capacity for visual discrimination.

McLaughlin agrees that counterfactuals of the sort to which Lewis adverts are properly associated with the exercise of a capacity for visual discrimination. (McLaughlin 1996: 205) This explains their attraction to the theorist. Normally, when the capacity is

exercised the counterfactuals obtain. However, McLaughlin continues, they do not *always* hold (as the counterexamples show) and are thus not constitutive of it.

Accordingly, on McLaughlin's account, we should explain what constitutes an experience as a *perceptual* experience as follows:

If a subject has visual experience and in having it exercises the capacity to discriminate a suitable range of scenes, then the subject sees;

if it is not the case that a subject has a visual experience such that in having it the subject exercises the capacity to discriminate a suitable range of scenes, ...then the subject does not see. (McLaughlin 1996: 224)

McLaughlin retains what he takes to be the kernel of truth in Lewis, while rejecting the talk of counterfactuals.

Part of the license for McLaughlin's view comes from the counterexamples, which we have already discussed. If my treatment of them is good, then he cannot receive any support for his account from them. In addition to the counterexamples, McLaughlin's position rests on a general opposition to the use of counterfactuals in giving constitutive accounts of the exercise of capacities. I'll discuss two lines of argument that converge on the denial that the exercise of capacities is constituted by counterfactuals: (i) the *Metaphysical Objection*, and (ii) the *Objection from Success*.

Let me turn to Colin McGinn for a statement of the Metaphysical Objection.

McGinn writes:

[I]t is unsatisfactory to employ counterfactuals in a primitive way in one's analysis of categorical propositions; they have dependent truth value. We can always legitimately ask what **makes** a given counterfactual true and expect to be presented with a suitable categorical fact. Now it seems to be that this general thesis imposes a constraint upon philosophical analyses, to the effect that we should be able to say what categorical propositions ground the counterfactuals we employ in the analysis....if non-circular categorical grounds can be produced it seems that the counterfactuals are in principle dispensable in the analysis; they

serve merely as an eliminable intermediate or interim step to the real analysis, which is categorical in form. (McGinn 1999: 16)

The point is something like this: Since counterfactuals are grounded in categorical facts, that the relevant counterfactual dependence obtains may be explained by reference to the relevant categorical facts. Since the counterfactuals can be so explained, their presence in the analysis is eliminable. Why is this? The explicability of the counterfactuals by categorical facts makes it reasonable to suppose that the exercise of a capacity ought to be *identified* with the categorical facts that ground the counterfactuals that typically obtain as a result of its exercise. The counterfactuals are explanatorily eliminable because exercising a capacity is identical to the occurrence of events which ground the counterfactuals that typically obtain. Thus we get McLaughlin's position: seeing requires the exercise of a capacity for visual discrimination, but the exercise of this capacity should not be given a counterfactual account.

A similar line of argument proceeds by considering the explanatory relations between capacities and our successes. Consider: If I know how to do something, then I possess a kind of capacity – the capacity for succeeding in some kind of endeavor. Consequently, associated with my knowing how, for example, to drive a car is some pattern of counterfactual dependence between my endeavors to drive and my success in so doing. Suppose that I were an expert driver of cars and were called upon to drive a car in a snowstorm. Were I to succeed that success would be part of a pattern of counterfactual dependence – the pattern one that might be supposed to constitute knowing how. But, presumably, were I to succeed what would explain my success is that I employed my considerable know how. My knowing how to drive a car explains cases of success rather than vice-versa.

Likewise, McLaughlin urges, when I succeed in having experience that correctly represent the facts, we should explain this by reference to my having exercised a capacity for visual discrimination. This seems entirely correct. Now: why does this fact cause trouble for the acceptability of a counterfactual account of the exercise of the perceptual capacity?

Since your exercise of the relevant capacity explains cases of success (both merely possible and actual), cases of success cannot constitute your exercise of the capacity. That is: that you succeed in visual discrimination in nearby worlds cannot constitute your exercise of a capacity. Consequently, one's exercise of the capacity to see cannot be constituted by a pattern of counterfactual dependence. Successes are doing double duty here: both explaining and being explained by the exercise of capacities.

Thus we have two arguments for the claim that *if* seeing is constituted by the exercise of a capacity for visual discrimination, then it should be denied that seeing is constituted by the obtaining of counterfactuals.

2.2: ASSESSING THE MCLAUGHLIN-STYLE OBJECTIONS

Before I assess the two objections, I'd like to argue that we may accept their general conclusions, but resist their application to the case of perceptual capacities. Even if we grant McLaughlin's denial that capacities should receive a counterfactual account, it is still plausible that a Lewis-style account is correct. The reason is this: even if capacities do not receive a counterfactual account, certain capacities are capacities to do things that consist partly in counterfactuals.

Let me explain. Consider, for example, a machine's capacity to correctly represent the temperature. Just insofar as this is a capacity and we have agreed that capacities are not constituted by counterfactuals, we have no reason to think that the exercise of *this* capacity involves the obtaining of counterfactuals.

To succeed in correctly representing the temperature, it is enough if (in the actual course of events) the machine represents the temperature as 90 degrees when it is 90 degrees. However the capacity for visual discrimination looks different. A capacity for visual discrimination is plausibly counterfactual-involving insofar as *discrimination* is plausibly counterfactual-involving. In this respect, the capacity for visual discrimination differs from the capacity for correctly representing. These capacities differ because the kinds of success they underwrite are different.

Discrimination consists, roughly, in treating different things differently. However, one can succeed in one-off discriminations. For example, one can imagine a creature that is built to discriminate between red things and blue things. However, due to an impoverished environment, the creature only exercises its capacity once. One happy day, it encounters a red thing and, in virtue of its capacity for discrimination, responds to it correctly as red-and-not-blue. It has discriminated, but it has not actually treated the different relevant things differently. That is: it has not actually treated some red things as red-and-not-blue and some blue things as blue-and-not-red. Its one-off discrimination involved treating only *one* thing one way. What then might its success in discrimination consist in? If actual encounters cannot do the work of making it true that, roughly speaking, the creature treats different things differently, then perhaps merely possible encounters will do. You might think, therefore, that when the creature successfully

exercises its capacity to discriminate, it succeeds in discriminating, which is to say it succeeds in treating different things differently in non-actual situations. The fact that one-off discriminations are possible suggests that a counterfactual account of discrimination is warranted.

If a counterfactual account of discrimination is warranted, then, while capacities *as such* might not implicate counterfactuals, capacities for discrimination may well do so. Lewis is not clear on the relation between his counterfactual theory of perception and the claim that perception involves the exercise of a capacity for discrimination. Did he mean his analysis to follow from a general reduction of capacities to counterfactuals or did he mean his analysis to apply as a result of the kind of capacity perception involves? I am not sure. At any rate, we can protect a Lewis-style account against a general denial of the reducibility of capacities to counterfactuals by adverting to the kind of thing discrimination is.

I have been assuming that capacities do not reduce to counterfactuals, but is this assumption warranted? What are we to say about the Metaphysical Objection and the Objection from Success?

The Metaphysical Objection proceeds from the claim that counterfactuals are grounded in categorical facts to the claim that counterfactuals might be eliminated from an analysis in favor of their categorical grounds. I am sympathetic to the premise of this objection – i.e. that counterfactuals are grounded in categorical facts. However, it is not clear to me that this establishes either that the exercise of capacity to see ought to be *identified* with the grounds of the associated counterfactuals or that the counterfactuals form an *eliminable* part of the analysis.

Everyone in this debate agrees that counterfactuals are associated with the exercise of capacities for visual discrimination. Presumably the associated counterfactuals have categorical grounds. However, it is further presumable that the categorical grounds for these counterfactuals are wildly disjunctive. In this respect, the relation between counterfactuals and their grounds is like the relation between mental states and the physical states on which they supervene. What grounds the counterfactuals associated with the exercise of my perceptual capacities can be expected to differ from what grounds those associated with the exercise of, say, an alien's perceptual capacities. In other words, counterfactuals are *multiply groundable*. In one case, the counterfactuals are grounded in *this* way. In another case, the counterfactuals are grounded in *that* way.

Thus, though it might be possible in *particular* cases to explain things without reference to the relevant counterfactuals, such explanations will have to proceed piecemeal and so fail to unify phenomena that would otherwise be unified. Suppose, for example, you think that a certain kind of epistemic value is constituted by a belief's being caused by an episode of seeing. You might wonder, regarding a particular belief, why does it count as possessing this kind of epistemic value? If counterfactuals are eliminated from the account in favor of their categorical grounds, then the explanation will take the following form: this belief was the upshot of such-and-such causal pathway. But such an account can be expected to leave out, given the fact that counterfactuals are multiply groundable, what such-and-such causal pathway has in common with the other causal pathways that constitute the epistemic value.

Thus elimination of counterfactuals in favor of their categorical grounds can be expected to fail to unify phenomena that could otherwise be unified. For this reason we should reject the Metaphysical Objection.

What about the Objection from Success? This argument involved a kind of slight of hand. The basic worry was something like this: on counterfactual accounts of the exercise of a capacity, explanations of success are circular explanations. Circular explanations are problematic, of course. For example, it would be a problem if the causal explanation for P invoked Q – where Q is supposed to be causally explained in turn by P. We'd have a kind of circular account of how P happened.

However, the explanations at issue in the Objection from Success are different kinds of explanations. We might *causally* explain someone's success in a venture by citing their capacities, but, according to counterfactual accounts, cases of counterfactual success constitute the exercise of the capacity. Thus, we *constitutively* explain an occurrence as the exercise of a capacity by noting the counterfactuals. Such-and-such is constituted as the exercise of a capacity because it occurs as part of a pattern of success. We do not thereby causally explain that occurrence.

It is perfectly compatible with this kind of constitutive explanation that the possible successes that constitute an occurrence as the exercise of a capacity are themselves causally explicable by that capacity. There are fairly non-controversial cases in which effects of a cause constitute that cause as such-and-such. For example: Imagine that pushing a button causes a war. You push the button and so start the war. The fact that the war resulted from your pushing of the button *constitutes* your button-pushing as, also, a war-starting. (Though not necessarily an intentional war-starting.) In this case,

we constitutively explain something in virtue of its effects. I see no reason not to do the same in the case of capacities.

The metaphysical objection and the objection from success are both unsuccessful justifications for denying that counterfactuals might constitute something as the exercise of a capacity for visual discrimination.

Suppose that McLaughlin is correct and that the exercise of a capacity for visual discrimination is *not* constituted by counterfactuals: what room is there for counterfactuals in an account of seeing? At this point it is helpful to remember that McLaughlin is giving an account of *intransitive* seeing. I am interested in *transitive* seeing. Thus our concerns are to a certain extent distinct.

Even if McLaughlin is correct and the exercise of a capacity for visual discrimination is not constituted by counterfactuals there might still be room for counterfactuals in an account of transitive seeing. The reason is this: That an experience is the upshot of the exercise of a capacity for visual discrimination does not yet tell us what its objects are – since a seeing of one thing might not be a seeing of another, even as both are *seeings*.

The selection problem and the problem of causal deviance remain as motivations for accepting a counterfactual account of transitive seeing. What could McLaughlin's account say about them? It should be obvious that McLaughlin's theory of intransitive seeing provides no hints about how it might be extended to solve the selection problem. You might suspect, however, that McLaughlin's account provides a solution to the problem of causal deviance. It would go something like this: what makes a causal chain ending in an experience *deviant* – and so insufficient for seeing a thing – is that causal chain fails to add up to an exercise of a capacity for visual discrimination.

This would be a mistake. Often enough deviance comes from interventions in the functioning of the eyes or the brain. This kind of deviance might be ruled out as a case of seeing because such causal chains do not add up to the exercise of a capacity for visual experience. In the raven case, the scientist's panic is a source of deviance only because it alters the normal functioning of the visual system. Maybe the reason you don't see the raven, despite having experience caused by it is that your raven-y experience is not the expression of a capacity for visual discrimination. Adverting to the exercise of capacities seems like it could account for the raven case.

However, we can imagine cases of cornea-external deviance – i.e. cases in which the causal chain leading from the eyes inward is in order, but in which the causal chain leading from the surfaces of objects to your cornea is deviant. Consider the case of the

Censorious Window:

In the future, windows will be designed to allow through only such light as would protect our privacy. Should someone inside a building be visible through such a window, the window would bend the offending light passing through it to transmit an image portraying some random scene. Now it happens you are walking by a house and glance into a censorious window. There is a family eating dinner inside and so the censorious window bends light to portray some random scene. Accidentally, the scene portrayed is qualitatively identical to the scene inside the house. It appears as if there is a family eating dinner *just so*. And indeed there is a family eating dinner *just so*.

The family has caused your experience, but nonetheless you do not see them. Why?

There was no interruption to your retina or brain. Your visual system works as it should.

It seems to me that there is no reason at all to think that the experience you have as you peer into the window is not the upshot of an exercise of a capacity for visual discrimination. Thus, while we might hope to use McLaughlin's position to eliminate

cases of causal deviance, we shall have to advert to something else. How about counterfactual dependence?

What I have done in this part of the chapter is justify and defend a refinement and extension of Lewis' counterfactual analysis of intransitive seeing. I have diffused some of McLaughlin's general grounds for doubt about the applicability of counterfactuals in the analysis of visual perception.

PART 3: PHOTOGRAPHS AND ROBOTS

3.1: THE CASE OF TIM AND TOM

In his article, "A Causal Analysis of Seeing", Michael Tye provides a counterexample that sheds doubt on the claim that counterfactual dependence is sufficient for perceiving a particular.²⁶ Like me, Tye accepts a counterfactual account, but thinks it needs to be supplemented. However, his purported counterexample is not in fact a counterexample. As I'll argue, we can look to photographic representation to see how Tye has gone wrong. Since we can see a thing by seeing a photograph of it, Tye's counterexample does not succeed.

Tye argues that counterfactual dependence is insufficient for seeing a thing. I paraphrase Tye's counterexample as follows:

Tim and Tom: There are two robots, Tim and Tom. The robots have been designed such that Tim controls Tom by remote control. Anything motion Tim makes Tom makes. Tom is located right in front of you in a well-lit room. Tim is located in an identical room not far away. (Tye 1982: 320-21)

²⁶ Tye does not present his analysis as a refinement of a counterfactual theory, speaking instead in terms of systematic variation. However, if there are differences between a tracking relationship and systematic variation, they do not matter here. Accordingly, I will construe Tye's proposal as a proposal about a counterfactual theory of perception.

The problem is that in this case your experience counterfactually depends on both Tim and Tom. Changes in each determine changes in your experience. Moreover, the dependence is robust: we may suppose that the connection between Tim and Tom is not fragile and so, in all nearby possible worlds, your experience witnesses a tracking relation with both Tim and Tom. However, it is plausible that while you see Tom, you do not see Tim. Thus, Tye urges, a counterfactual dependence analysis needs supplementation.

3.2: A CONCESSIVE RESPONSE: TYE'S PROPOSAL

Why, according to Tye do you see Tom, but not Tim? Let me present an extended quotation to represent Tye's reasoning:

Originally, Tim's role in producing my experience is merely one of causing Tom to take up his position in front of the mirror. But as Tim's M-properties are systematically varied, Tim's causal role becomes much more complicated in virtue of his relationship to Tom taking on the new programmed connections which transmit changes in M-properties from the former robot to the latter. Now clearly if we were to hold *fixed* the various programmed factors which are responsible for Tim's changing his causal relationship to Tom in this way (while leaving other factors alone), the spatial properties of my robotlike sensum would not systematically vary as Tim's M-properties are varied. ... But I do see Tom, since the manner in which Tom brings about my visual experience remains substantially the same whether or not his M-properties are varied. (Tye 1982: 322)

[T]he M-properties of a material object x are those spatial properties of x which have counterparts in the spatial properties of the experienced visual sensum. So, for example, in the case of the man viewing a distant star, shape is *not* one of the star's M-properties. (Tye 1982: 319)

On Tye's account, you don't see Tim because situations in which the features of Tim that might be perceived are varied are also situations in which his causal role in producing your experience varies significantly. You see Tom because in situations in which his

perceivable features are varied, Tom's role in producing your experience stays substantially the same.

Tye's account is like mine in requiring that states of experience track changes in the seen object, holding fixed the process connecting one's experience and the object of perception. However, it is only in *some* cases, according to Tye, that we should hold fixed the process connecting one's experience to the object of perception. If (as with Tim) the nature of the actual process would change substantially, were the perceptible features of a candidate object of perception to change, then we should hold the actual process fixed and ask whether the tracking relation obtains. By contrast, if (as with Tom) the nature of the actual process would not change substantially and were the perceptible features of a candidate object to change, then we should allow the process to vary and ask whether the tracking relation obtains.

On Tye's account, then, were Tim's position were different, his role in producing an experience would be substantially different, and so we must hold the process linking him to Tom fixed and ask whether a tracking relation obtains between states of Tim and your experience. Since the actual causal path is such that states of Tim only affect your experience by affecting states of Tom, Tye concludes that holding the actual process fixed would not allow for the obtaining of a perception-grounding tracking relation. As Tye puts it:

[C]learly if we were to hold *fixed* the various programmed factors which are responsible for Tim's changing his causal relationship to Tom in this way (while leaving other factors alone), the spatial properties of my robotlike sensum would not systematically vary as Tim's M-properties are varied. (Tye 1982: 322)

Held fixed, according to Tye, the process linking Tim to Tom is such that it fails to change Tom to reflect Tim's states. Tye seems to be thinking of holding processes fixed

as holding fixed the causal chain emitting from Tim and terminating in Tom's bodily position. This would prevent changes in Tim from being mirrored in Tom – and consequently your experience.

I am suspicious, however. It is true that changes in Tim's position would bring about changes in the causal role he plays in producing your experiences. Are these substantial changes? I am not sure. So long as the mechanisms in virtue of which Tim generates and transmits information about his position to Tom are doing what they ought to, then it seems to me that the process connecting Tim and Tom's positions is intact and Tim's causal role is not significantly different.

Indeed, it seems to me that even if we accept that the relevant counterfactuals are one in which we are to hold the process linking Tim and Tom fixed, we'll reach the conclusion that a tracking relation obtains between your experience and Tim. Holding fixed the process would require, for example, allowing the sensors in Tim's body to transmit their representations of Tim's body to that part of Tim that generates commands for Tom based on them. By contrast, Tye must be supposing that holding the process fixed means holding fixed the content of the signals passing from Tim to Tom such that even when Tim is in a position that differs from his actual position he still transmits the same signal as in the actual case.

So I do not think Tye's treatment of the case of Tim and Tom is successful. However, if you are persuaded that you cannot see Tim, it might be possible to advert to the details of the information channel that connects your experience and states of Tim to explain why you cannot see Tim. Maybe, for example, you would reason as follows: The organism-external channel in virtue of which the tracking relation obtains must transmit

information via light. Since the information reaching you about Tim's states is not transmitted entirely via light, you do not see Tim.

The fact that it is a scientific discovery that visual information is transmitted by light does not rule out this possibility. However, I think it is possible to see, even when the organism-external information is not carried by light. Consider the movie *Daredevil*. In this movie, a man who has lost the use of his eyes develops amazingly acute hearing, such that he is able to rely on his hearing to gather much of the same spatial information ordinary perceivers gather by the use of their eyes. The movie represents his subjective state by the use of monochromatic images. The suggestion of the movie, then, is that his visual cortex has become connected to his ears in such a way that information coming from his ears generates visual experiences. Thus, while it is as if he is blind (since his eyes are nonfunctioning), it is also as if he can see (since he is subject to visual experiences that track changes in his environment). He can see via information carried by sound. I think one can imagine many other variations on this case. Perhaps, for example, a very large creature could see in virtue of being endowed with organs that detect changes in the amplitude and wavelength of gravitational radiation. It would be a mistake, therefore, to require that seeing exploit only information carried by light.

3.3: PHOTOGRAPHIC REALISM AND SEEING TIM

While it is initially plausible that we do not see Tim, I have come to the view that we *can* see Tim. Thus we need not supplement our counterfactual theory with restrictions designed to rule out seeing Tim. One reason for this view would be that Tim satisfies the conditions on an object of perception laid out in a true theory of the

perceiving relation. Of course, in this context this reason is insufficient. That you see Tim is made rational both by the independent strength of the counterfactual theory of perception and by considering the analogies between Tom and a photograph.

Ken Walton has argued persuasively that you can see a thing by seeing a photograph of it. As it turns out, the relationship between Tom and Tim is like the relationship between a photograph and what it pictures. Thus, if you can see the object of a photograph by seeing a photograph, then you should be able to see Tim by seeing Tom. In this section, I will present Walton's view and use it argue that Tye's attempted counterexample is not actually a threat to the counterfactual theory.

3.4: PHOTOGRAPHIC REALISM

Ken Walton argues that in looking at a photograph it is possible to see what it is a photograph of. (Walton 1984) Walton's claim is justified as part of an inference to the best explanation for the judgment that photographs are somehow more realistic than even the most realistic paintings or drawings.²⁷

The superior realism of photography, whatever it should turn out to be, shows up in many of our practices. As Walton puts it:

Photographs of a crime are more likely to be admitted as evidence in court than paintings or drawings are. Some courts allow reporters to sketch their proceedings but not to photograph them. Photographs are more useful for extortion; a sketch of Mr. X in bed with Mrs. Y – even a full color oil painting – would cause little consternation. Photographic pornography is more potent than the painted variety. Published photographs of disaster victims or the private lives of public figures understandably provoke charges of invasion of privacy; similar complaints against the publication of drawings or paintings have less credibility. I expect that most of us will acknowledge that, in general, photographs and paintings (and comparable non-photographic pictures) affect us very differently. (Walton 1984: 246-47)

²⁷ In what follows, I'll use paintings as representative non-photographic pictorial representations.

We might understand the truth of the claim that photographic pictures are supremely realistic to simply consist in these kinds of facts about our treatment of photographs. What, then, would explain and rationalize the differences in our treatment of photographs and paintings?

Walton argues that the explanation cannot be that photos are invariably more informative representations of their objects. Some paintings – the so-called photorealistic ones – are indistinguishable from photographs taken of the scenes they picture. Some photographs – those that are, for some examples, produced using distorting lenses, digital manipulation, or dark-room tricks – are less truthful than a painting of the same scene might have been.

On Walton's account, what explains the thought that photographic representation is the *sine qua non* of pictorial realism is the fact that one can see the object of a photographic representation by looking at it, while one cannot see the object of a painting by looking at it. Just as seeing a thing might be less informative than having it described – if one's glimpse of it is, for example, quick and partial, photographs needn't be particularly informative to receive the special treatment we reserve for them: "We can't expect to acquire any particularly important information by looking at photographs [of departed loved ones] which we have studied many times before. But we can *see* our loved ones again, and *that* is important to us." (Walton 1984: 253)

What would explain why you can see the object of a photograph by looking at the photograph, while you cannot see the object of a painting by looking at it? Walton argues that the crucial difference between photographs and paintings lies in their history. In particular, a photographic representation is mechanical in a way that a painting is not.

Unlike a photograph, a painting's representation of a scene is mediated by its creator's belief-like attitudes. Or, at any rate, since photographs don't take themselves, the role belief-like attitudes play in creating photos differs from the role they play in creating paintings. Walton cashes out the difference in counterfactual terms:

Photographs are counterfactually dependent on the photographed scene even if the beliefs (and other intentional attitudes) of the photographer are held fixed. Paintings which have a counterfactual dependence on the scene portrayed lose it when the beliefs (and other intentional attitudes) of the painter are held fixed. (Walton 1984: 264)

Thus, that a picture is a candidate for supporting perception of the pictured thing depends on the counterfactual dependence between the contents of the picture and the states of the pictured scene. At its core, Walton's account of the fact that one sees what is photographed by seeing a photograph of it depends on the historical similarity between a perceptual experience and a photograph. Since a perceptual experience counterfactually depends on its objects, and since a photograph counterfactually depends on its objects, one can see a thing by seeing a photograph of it.

However, according to Walton, in order that this dependence support perception, it is necessary that the scene-depiction dependence not be wholly dependent on propositional attitudes. That is, it cannot be strongly mediated by propositional attitudes. In the case of paintings, the attitudes of the creator make all the difference, while in the case of photographs they make no difference to whether a tracking relation obtains.

Walton justifies the claim that strong mediation by propositional attitudes rules out perception by invoking the case of Helen's benevolent neuroscientist: The neuroscientist wishes that Helen – who has lost her eyesight, but gained a prosthetic experience-producer – enjoy experiences that match the scene. He manipulates the

prosthetic experience-producer to generate experiences in Helen that match the scene he thinks he would see were he in her position. Suppose that he is both very good at this and extremely dedicated. Thus, in virtue of the neuroscientist's benevolent interventions, a robust counterfactual tracking relation obtains between Helen's experience and the scenes in front of her. Walton claims that while it seems to Helen as if she sees, she does not. (Walton 1984: 265)

Now, as to why strong mediation by propositional attitudes might make the difference between seeing and not seeing, Walton does not explain. This is a pressing question because in the case of Helen and the benevolent neuroscientist we have reason to doubt either the sufficiency of a counterfactual theory of perception or the judgment that Helen does not see. Additionally, we have reason to believe that the beliefs of a perceiver make some difference to the features represented in experience. So one might expect that even in typical cases of human perception a world-experience match might be weakly mediated by propositional attitudes. Zenon Pylyshyn argues that while early visual processing proceeds independently of our belief-like attitudes (except for attention-focusing effects), higher levels of visual processing do not. (Pylyshyn 2003: Ch 2) In particular, *seeing as* – the process that (for example) distinguishes the duckish experience from the rabbitish experience of a duck-rabbit image – is sensitive to our beliefs.²⁸ Thus, for example, if I display the duck-rabbit to you and say “check out this duck”, you're likely to be subject to duckish experience.²⁹ Thus, unless we are supposed to be quite

²⁸ It is important to realize that *seeing as* and *seeing that* are different kinds of mental states. In particular, *seeing that* is a species of judgment, while I have just argued that *seeing as* is not. As far as I can tell, when we say that someone can see that *p*, what we mean is that they judge that *p* on the immediate basis of their perceptual state. Moreover, seeing that *p* is probably a species of knowledgeable judgment. If you can see that *p*, then you know that *p*.

²⁹ It is an open possibility of the defender of the hypothesis that seeing cannot be mediated by belief-like attitudes to deny that *seeing as* is strictly speaking a form of seeing. However, this seems like gratuitous

generally blind, then we cannot suppose that simple mediation by belief is a barrier to seeing.

So why would strong mediation be a barrier to perception. Luckily we need not settle this question. The reason is this: I am suspicious about our judgments in cases like the case of Helen and the neuroscientist. In order that the neuroscientist's intervention might support a robust pattern of counterfactual dependence, we must suppose that the neuroscientist is *very good* at determining what kind of experience Helen should have and monomaniacally dedicated to bringing it about that she has this experience. It might be that in imagining the case of Helen and the neuroscientist we are unable to suspend our default assumption that the neuroscientist, like other agents, has an open future. That is, insofar as we suppose that the neuroscientist is a person, we suppose that his interventions are a matter of decision. When he thinks some experience would be appropriate for Helen, it is a matter of decision whether to make it true that Helen is subject to that experience. Thus, in nearby possible worlds, there are some in which he provides her with the experience and others in which he doesn't. Thus, insofar as we tacitly suppose he is an agent, we'll find it hard to genuinely suppose that the neuroscientist's interventions are sufficient for grounding the right counterfactuals.

Thus volitional attitudes cause trouble for seeing via paintings or neuroscientists. We conceive of volitional attitudes as indeterministic and so, if a representation is mediated by them, we cannot suppose that a tracking relation obtains between the

linguistic stipulation to me. For example, while seeing as might be sensitive to our beliefs, it is phenomenologically more like seeing than judgment. I do not feel as if I am *concluding* that that is a duck. If it were a judgment I was making as I gazed at the duck, I would expect that my experience would be partly constituted by an inwardly asserted sentence – something like “that’s a duck.” The difference between my duckish experience and my rabbitish experience is not a difference in the words I say to myself. It comes on with apparent automaticity. It feels like a specifically visual difference – even if it is a difference that can be influenced by belief-like states. Since the phenomenology of seeing-as is vision-like, it seems gratuitously strict to eliminate seeing-as from the category of genuine seeing.

representation and states of represented scene. If our conception is true, then in actuality no one could see by a process that is strongly mediated by volitional attitudes.

My hunch is that Walton has made a mistake in supposing that photographic and painted pictorial representations *must* differ with respect to whether one can see what they picture. He is, however, probably correct that in *our* world they differ. So long as the benevolent neuroscientist truly is able to make the relevant counterfactuals obtain, then Helen sees. But no actual human neuroscientist can do this. Likewise, if a painting depended on the states of what it pictures as a photograph does, then it would enable indirect seeing. However, no actual painting is likely to be like this – given the involvement of agency in determining the content of a painting.

3.5: SEEING TIM

But now note: the relationship a time slice of Tom bears to Tim is analogous to the relationship a photograph bears to what it is a photograph of. Even better, the relation Tom bears to Tim is analogous to the relationship a video image bears to what it pictures. Thus, if we see a thing by seeing a photograph of it, we can see Tim by seeing Tom.

If we see Tim, we see him *indirectly*. Thus we see Tim by seeing Tom in the same way we see a person by seeing part of him. Our readiness to say that we have seen a thing, when we have seen it indirectly, is sensitive to context. Indeed the context dependence is such that people will easily swing back and forth between affirming and denying that an indirectly seen thing has been seen. The swings depend on what kind of importance is placed on certain features present in ordinary cases of perception, but lacking in cases of indirect perception. Thus, for example, when I see a thing by seeing

part of it, I might be in no position to gather relevant information about it. Emphasize this aspect of things and folks conclude that the indirectly seen thing was actually unseen.

The barrier to saying that we see Tim is that Tim is in a different room and so occluded by walls and so forth. Thus Tim differs from ordinarily seen things. Of course, if you can see a thing by seeing, for example, a video of it, then occlusion is not invariably an obstacle to perception. A security guard might keep watch on the grounds of a factory by viewing live video of it. The objects revealed to him might be occluded by walls and, indeed, by the video screen at which he gazes, but nonetheless he sees them. Since Tim is like the photographed thing, I think we should say that we see him. Occlusion matters to perception only contingently.

Perhaps it would be helpful to imagine that instead of discovering photography, our civilization had developed transparent boxes containing a malleable substance that could be formed by automatic processes into accurately colored, small-scale models of mid-sized dry goods. Call these *fmodels*. Imagine that the correspondence between the *fmodels* and their representational targets typically is (or can be) very good. That is, suppose that there typically is a tracking relation between *fmodels* and their representation targets that corresponds to the tracking relation between photographs and their objects.

With sufficient technological development, these boxes could function like televisions, providing moving tableaux. So, instead of watching televised animal documentaries, for example, you'd watch *fmodels* of exotic animals doing interesting things. Instead of surveying the grounds using video, a security guard would watch for suspicious changes in his territory, by looking at the evolution of his *fmodels*. You might

collect still fmodels – or the computer files that could construct such fmodels – of important people and events in your life. Instead of comparing paintings and photographs, Walton would have compared fmodels and statues, advocating for the transparency of fmodels. These practices would not be irrational. That is, since fmodels are like photographs, the practice of treating them as such would make sense. Now what is Tom, but a fmodel of Tim?

PART 4 THE THEORETICAL POSSIBILITY OF ELIMINATING REFERENCE TO CAUSATION

The theory I have given is a version of a causal theory. That is, it invokes causation as part of its account of the seeing relation, even as it supplements causal dependence with counterfactual dependence. However, it is possible that seeing doesn't require causal dependence and that we need only advert to counterfactual dependence. The possibility I have in mind comes from Malebranche via Dummett: "If someone believes, with Malebranche, that the presence of the object and my experience of it are joint effects of some further cause, his belief does not violate the concept of perception, so long as he allows that my perception supplies a reason for taking the object to be there." (Dummett 1979: 35-36) There are two suggestions here. First, it is essential to perception that to perceive is to have a reason for taking an object to exist hereabouts. Second, it is not essential that to perceive there must be a causal relation between one's experience and the objects of the episode of perception of which it is a part.

It is not clear that perception is essentially a reason-provider and so it is not clear that Dummett has any reason for accepting the second suggestion. However, even if we

think that perception need not be a reason-provider, we can imagine cases in which an experience counterfactually depends on an object, but without causally depending on it. I have in mind, of course, cases in which states of experience and states of the candidate object of perception do not causally interact, but are both effects of a common cause (e.g. ideas in the mind of God). To the extent we see in such cases, we'll have overriding reason to eliminate reference to causation in our analysis. Now, since I am most interested to defend the counterfactual account, I have no great anxiety about whether, in the end, reference to causation is eliminable.

There is one reason that reference to causation might be implicit in my account. It is epistemically possible that it is necessarily the case that a process takes place if and only if causation occurs. If this epistemic possibility is actual, then insofar as my account makes reference to processes, it is implicitly committed to the claim that any case of perception is also a case of causation. As we will see below, a process is a puzzling thing and it is not clear how we should decide this question. Presumably, answering it will involve some philosophical trench warfare in which we think about sequences of events about which we will have to simultaneously decide whether causation is occurring and whether such sequences constitute the realization of a process. These cases will have to be fairly *recherché* since every sequence of actual events that I can think of that clearly constitutes the realization of a process is also a sequence of causally-related events.

However, I do not need to decide it in this section. It might be that any counterfactual account that, like mine, invokes processes is implicitly committed to the claim that if perception is occurring, then causation is involved. It suffices to return our attention to the Malebranche case to see how the role causation would be playing in such

an account differs quite a bit from the role it plays in causal theories of perception. Causal theories of perception characteristically claim that there is some causal relationship between objects of perception and experiences in virtue of which those objects are objects of perception. A counterfactual account that accommodates intuitions about Malebranche cases and also includes reference to processes might be implicitly committed to the claim that perception is causation-involving, but notice that the causation it posts will not necessarily flow from the objects of perception to experiences. The Malebranche cases are cases of counterfactual dependence secured by the joint and branching cause of states of an object and states of experience. Thus, a counterfactual account that accommodates the Malebranche cases might still assert that perception is causation-involving (insofar as processes are supposed to be causation-involving), but not assert that the process (and so the causal sequence) that grounds the counterfactual dependence is one that begins with the object of perception.

PART 5 THE GENERAL GENERALITY PROBLEM AND ME

It is arguable that if seeing requires that actual processes be held fixed, then we see nothing. The problem is that, on a certain construal of what it would be to hold actual processes fixed, a tracking relation will not obtain between experiences and states of the environment. Suppose that holding the actual process fixed entailed not allowing any variation of any intrinsic feature of the causal chain actually connecting one's experience to object in the environment. For example, suppose that in the actual case, an object

reflects light of a certain wavelength and intensity – let’s say the wavelength is such as to make it appear blue. This makes it true that you are subject to an experience that represents a blue thing with such-and-such extension located at such-and-such location. Now if holding fixed the actual process rules out allowing any variation of any feature of the actual causal chain, then in nearby worlds in which the object would reflect a different wavelength – one such as to make it appear red – the process would still eventuate in an experience that represents it as blue. Likewise, other variations in its color, location, or extent would not be matched by the resulting experience. My theory, the worry goes, has the consequence that we see nothing. This would, I concede, be a dire problem for my theory.

The problem here is similar to the generality problem raised by Conee and Feldman for reliabilist theories of epistemic warrant. (Feldman 1995) (Conee & Feldman 1998) Consider a case of belief that *p*. According to a process reliabilist, this belief that *p* is warranted iff it is the issuance of a belief-forming process that reliably produces true beliefs. The generality problem arises because the token process that produced my belief that *p* will always be an instance of many types, for example those characterized by the following principles:

- Believe *p*, given *e* (where *e* is some evidence that you have).
- Believe *p*, if you are having visual experiences of with content *c*.
- Believe *p*, if you are in Ann Arbor.
- Believe *p*, if *p*.

The problem is that these types vary greatly in their reliability. A good case, one that ought to count as warranted belief, might be counted as unreliable (and so unwarranted)

on some ways of understanding the process that generated it. A bad case, one that ought to count as a case of unwarranted belief, might be counted as reliable (and so warranted) on other ways of understanding the process that generated it. Until we have a principled way of picking out ‘the’ process that formed my belief, we have no principled way of evaluating its warrant. The reliabilist owes us an account of what settles this question – an account of what makes it the case that one process, and not the others, is the relevant process to evaluate in evaluating the epistemic quality of my belief. The problem constitutes an objection to reliabilist theories of epistemic warrant, once one becomes suspicious that no solution is in the offing.³⁰

Similarly in the case of perception, some ways of understanding which process generated an experience will yield the incorrect results – the result that something in fact seen was unseen, the result that something in fact unseen was seen. So, in requiring that the actual process be held fixed in evaluating whether a tracking relation obtains, what am I requiring? It is not my intention here to give a complete answer to this question, since I do not have one. Instead, I have the more modest, and more attainable, goal of pointing the way toward an answer, delineating some constraints on what an answer might be.

It is also worth noting that the problem here is not mine alone – nor is it shared only with reliabilists about epistemic warrant. Ralph Wedgwood and Juan Comesaña have recently, and persuasively, argued that every reasonable epistemology – whether externalist or internalist about the possession of epistemic value – will ultimately have to

³⁰ It would be nice to see a principled argument that there could be no solution. As it stands, the pressure on the theorist who makes reference to processes is merely the pressure to be completist. I like a complete account as much as the next guy, though I have yet to see one and it seems unfair to require one.

resolve the generality problem. (Wedgwood 2002: 287) (Comesaña 2006) My theory is, therefore, in good company.

A process is a puzzling thing. A process can be both a sequence of events (and so concrete or located at a time) and something like a recipe (which is abstract and repeatable). The problem here stems from the fact that a process is something the occurrence of which can be referenced in answering a kind of ‘how’-question. (How did you come to live in Idaho? How did you bake that cake? How did the fundamental particle end up over there?) We can answer this kind of question by indicating the sequence of events at the end of which lies the target of our inquiry.³¹ However, until we know how a sequence of events constitutes a unity, it seems that we do not yet know what we wanted to know in wondering how something happened. This suggests that we should think that a process is a sequence of events constituted as a unity.

Moreover, the specific unity constituted by an actual sequence of events can be constituted by rather different sequences of events. Suppose I demonstrate to you how to bake a cake. Having finished baking the cake, I tell you, “Now do the same thing.” I have commanded you to try and replicate the process by which I baked the cake. However, it is compatible with your satisfying my command that the sequence of events leading up (hopefully) to your cake differs in many interesting ways from the sequence of events that led up to my cake.³² For example, in measuring the flour I might have used my right hand to manipulate the scale and you might have used your left hand. It could be that your sequence of events nevertheless instantiates the same process as my sequence. Or, perhaps, sensing a difference in the level of humidity in the air, you added

³¹ By ‘event’ I do not mean to refer to property instantiation. Rather, I mean to refer to time slices.

³² And not just with respect to its being you who were involved with the sequence.

less milk to your batter. There is a rather different sequence here, but the same process. Consider, for another example, two sequences that realize the process of using a radar gun to determine the speed of a vehicle. Presumably, in both cases the same fundamental physical laws are instantiated between pairs of events, but the events will instantiate these laws in different ways. However, in each case, the pattern of radio waves will differ with respect to their frequency and amplitude along different dimensions. And yet one might truly say that both sequences realize the same process – the process of determining the speed of a vehicle with a radar gun.

A given sequences of events might be truly described in ever so many ways and this is the root of the generality problem. An actual sequence satisfies a great many descriptions – to make thinking about this easier just consider three descriptions: D1, D2, D3. In one set of nearby worlds there are sequences that satisfy D1. In another, there are sequences that satisfy D2. In another, there are sequences that satisfy D3. The critic wants to know which of these worlds are worlds in which the actual process occurs. Or, assuming that the critic has some intuitive grasp on what kinds of changes to the actual situation constitute a change in process, the critic wants an explicit account of what we're doing when we hold the process fixed.

It is evident that in holding fixed the actual process it is insufficient to require merely that the relevant sequences satisfy *at least one* description that is also satisfied by the actual sequence of events. Not every such description will correctly identify the process – that is, not every such description will identify the relevant way in which the actual sequence of events is constituted as a unity. For example, that a sequence that realizes the process of baking a cake occurs on Tuesday, is presumably not a fact that has

anything to do with the way(s) in which it is constituted as a unity. Thus, if we want to know what would happen if that process were repeated, we would be gratuitously strict to consider only sequences that occur on Tuesdays. Moreover, and for the same reason, it would be a mistake to require that such sequences satisfy *every* description satisfied by the actual sequence.

What the critic demands is an account that tells us which of the descriptions the actual sequence satisfies need be satisfied as we hold the process fixed. Now some descriptions of a sequence of events are such that any sequence of events that satisfies that description constitutes the same kind of unity as the actual sequence. What the critic wants is an account that puts some conditions on these descriptions, some account that picks out the relevant descriptions.

It is one constraint, I claim, that these descriptions pick out natural process-kinds. At each world (among the worlds that provide the truth conditions for the relevant counterfactuals), the process-realizing sequence of events must instantiate the same natural process-kind.³³ Recall, for example, one of the principles in my list of principles according to which one might type the process by which I came to believe that *p*: Believe that *p*, given that you are in Ann Arbor. The sequence of events that eventuated in my believing that Lormand is amazingly smart is one that occurred in Ann Arbor and a description in those terms picks out the sequence. But, intuitively, that it occurred in Ann Arbor is an inessential feature of the process. The same process could have been spread

³³ Alston tries to use this kind requirement to solve the generality problem for reliabilist theories of epistemic warrant. (Alston 1995) Notice also that this proposal makes our judgments about whether perception has occurred ultimately hostage to science. We have some intuitive grasp on what kinds of things are not natural kinds – we can usually identify gerrymandered kinds as such. But if we are wrong about which sequences belong to which natural process-kinds, then we will be wrong sometimes about what the counterfactual effect of holding the process fixed would be. We are hostage to science, because it is an empirical matter what the natural kinds are.

out between Ann Arbor and Bangalore – thanks to telecommunications technology. This would explain why identifying a cognitive process as *believe that p in Ann Arbor* is problematic. It would be safe to assume that cognitive processes are not to be individuated by whether they occur in Ann Arbor or not. If we wish to cut nature at its joints, this will not do. Thus, if we want to evaluate the reliability of the process by which I came to believe that Lormand is amazingly smart, it would be a mistake to require that the sequence of events eventuating in my belief occur in Ann Arbor.

Now this might be enough to respond to the problem that started this part of the chapter. The worry, remember, was that a process-relative counterfactual theory of perception has the consequence that we see nothing. If holding fixed the actual process fixed requires holding fixed the microphysical description of the actual sequence of events, then in no case will changes to the object be tracked. The proper response to this might be as follows: this does not constitute holding the process fixed, since the microphysical description of the sequence that actually eventuated in your experience does not pick out a natural kind. At best, it picks out an arbitrary subset of a natural process-kind.

However, even if we can dissolve the puzzle with which I began, restricting the descriptions that pick out the relevant sequences to descriptions that pick out natural kinds, it is not enough to solve the larger problem toward which the puzzle gestured. Two sequences might instantiate the same natural process-kind, but intuitively differ with respect to which process they instantiate. Imagine that you are a scientist who works on the differences between the digestive systems of birds. Presumably, digestion is a natural process-kind. The subjects of your study (finches and penguins) are all digesters and so

the sequences of events in their events that you study instantiate the same natural kind, but you argue that the process by which a finch digests differs from the process by which a penguin digests. It could be, for a fictional example, that while finches (like cows) have an extra stomach that in which ingested food is fermented by micro-organisms, penguins do not.

Natural process-kinds may bear the determinate-determinable relation to each other. The processes by which the finch and penguin might digest differ from one another like a cerulean blue thing differs from a Prussian blue thing. They are both ways of digesting.

What about visual mechanisms? Does the same thing arise there as well?

Consider the case of the **Delicately Opaque Object**:

The delicately opaque object is opaque and colored blue, but were it differently located or shaped it would cease to be opaque and become perfectly transparent. Moreover, blue is the only color it could possess. Any change that would make it non-blue, would make it perfectly transparent. As it happens, the conditions are optimal for viewing it and you give it a good look.³⁴

It is plausible that you see the delicately opaque object. But does your experience of it track possible changes in it, holding fixed the actual process? Maybe we have a counterexample to process-relative counterfactual theories of perception. I agree that it is plausible to describe this as a case of seeing. If we held fixed merely that the sequence is one in which, for example, you open your eyes and receive information transmitted by light, then the tracking relation wouldn't obtain and we should have to conclude either that counterfactual accounts should be rejected or you don't see in the case of the Delicately Opaque Object.

³⁴ Thanks are due to Richard Grandy for suggesting this kind of case.

But this case is not, in fact, a counterexample. The circumstances in which the object becomes perfectly transparent are circumstances in which the process leading from the object to your experience is different. The difference is this: in the actual case, the object reflects light, while in the cases where the object is perfectly transparent, it transmits it. Holding fixed the process evidently involves holding fixed that the sequence is one in which reflection occurs (rather than transmission). That the light reaching your eyes from the location of the object is transmitted by the object rather than reflected by it makes a difference to which process has occurred. If we held fixed the fact that the process leading to experience is one in which reflection occurs, then, were the object differently colored, its different color would be matched.³⁵ Thus, we have a case in which holding the actual process fixed means holding fixed a determination of a determinable natural process-kind.³⁶

Maybe the case of the Delicately Opaque Object shows us that at each world (among the worlds that provide the truth conditions for the relevant counterfactuals), the process-realizing sequence of events must instantiate the same determination of a determinable natural process-kind.³⁷ However, this will not be enough. Determinables admit a hierarchy of determinations. Thus, while blue is a determination of colored, and cerulean blue is a determination of blue, there are shades of blue that are determinations of cerulean blue. Consequently, we should expect the same of the natural process-kinds an actual sequence of events might realize. As a result, we can expect some

³⁵ Notice that we'll be considering what would happen in worlds in which the delicate opacity of the object is not witnessed. It appears that, as Lewis thinks, the truth of counterfactuals can depend on worlds in which small miracles occur. Thus we needn't suppose that the nearby worlds – the truth-makers for counterfactual conditionals – need be physically possible.

³⁶ Notice also the role empirical knowledge is playing in the judgment that a change from reflection to perfect transmission is a change in process. It is only in virtue of my modest knowledge of optics and vision science that I am in a position to make this judgment.

³⁷ Provided, of course, that the determination is itself a natural process-kind.

philosophical puzzlement about whether two possible sequences realize the same process. One fix here would be to require that at each world (among the worlds that provide the truth conditions for the relevant counterfactuals), the process-realizing sequence of events must instantiate the same super-determination of a determinable natural process-kind that is also a natural-process kind. That is, the sequences must instantiate a natural process-kind that is a determination of a determinable natural process-kind that is not itself further determinable by a natural process-kind.

However, this fix won't work. Evidently, whatever description picks out the process will have to be one that picks out the *right* natural process-kind determination of a natural process-kind determinable. To see why this fix won't work we need to get a little clearer about the way in which processes are built from sub-processes.

A sequence might possess local unities. There are, therefore, sub-processes. However, holding fixed the actual process cannot simply be a matter of requiring that relevant possible sequences satisfy a description that picks out one among the many sub-processes realized by the actual sequence. For example, some part of the sequence of events that constitutes my making lasagna is constituted as a unity truly describable as *my boiling the pasta*. One might truly describe such a sequence as the process in which I boiled some pasta. This means that one cannot necessarily identify a process by truly describing some locally constituted unity. Making lasagna and making spaghetti are different processes, even if they both involve sequences that have regions that are unified insofar as they are pasta-boilings.

Not only would such a procedure misclassify processes, but it would leave my account open to a kind of censor case. If holding fixed the actual process is to prevent

sensor cases from being cases in which nothing is seen, we shall have to understand holding the actual process fixed as requiring more than that the relevant sequences satisfy a description that is true of some less-than-maximal region of the actual sequence. To see why this is so, consider the case of the **Almost Last Minute Censor**:

Imagine that the last part of the computation that eventuates in one's experience as one gazes at a white wall is the transmission of a signal, the content of which is something like *there is a white surface of such-and-such extent in such-and-such region*. Now imagine, as in the other censor cases, there is a censor who wishes that you be subject to whatever experience eventuates from such a signal. In worlds where the wall is different or nonexistent, the censor intervenes and produces the experience he wants you to have. The censor's intervention is late, but not last minute. What I mean is this: instead of directly producing in you the experience you might have as a result of looking at a white wall, he causes the transmission of the same signal that was transmitted in the actual case.³⁸

This would be a counterexample to a Lewis-style counterfactual theory – one which is not, like mine, process-relative. Is it also a counterexample to a process-relative theory? Is it the process in the situation in which the censor intervenes the same process as the process in the actual situation?

I think it is not a counterexample. The reason is this: it supposes that it is sufficient for holding the actual process fixed that we hold fixed only the last sub-process. That is, it supposes that the relevant sequences are ones that satisfy whatever description picks out the last region of the actual sequence. Some of these are sequences that realize the same process as the actual sequence (consider the sequences that are identical in every respect); others – for example, the sequences in which the censor intervenes – are not.

The problem here is not that description 'transmits a signal the content of which is something like *there is a white surface of such-and-such extent in such-and-such region*'

³⁸ Thanks are due to Eric Lormand for providing this example in conversation.

fails to pick out a natural process-kind of the relevant level of determination. We can suppose that this is as determinate and natural as process-kinds get. The problem is that it identifies the process by one of its sub-processes. Thus, holding the actual process fixed requires holding fixed the entire extent of the process. In line with the proposal that holding a process fixed (in assessing whether seeing has occurred) requires holding fixed a sequence's membership in a super-determinate natural process-kind we might propose the following: Each relevant sequence is such that, for each region of the actual sequence that realizes a super-determinate natural process-kind, there is a corresponding region that realizes the same super-determinate natural process kind.³⁹

There is something right about the requirement that each relevant sequence is such that, for each region of the actual sequence that realizes determinate natural process-kind, there is a corresponding region that realizes the same determinate natural process-kind. Sequences might belong to what I'll call *maximal natural process-kinds*. Two sequences belong to a maximal natural process-kind, iff each of their corresponding regions belong to the same natural process-kind and both possess such unity as a whole as to belong to the same natural process-kind. I think that it is just about right to say that in holding the process fixed, we're requiring that the relevant sequences belong to the same maximal natural process-kind.

The problem comes from requiring that they realize the same super-determinate natural process-kind. There are possible creatures for whom the kind of variations an experience must track – it is to be a perceptual experience – activate different super-determinations of a determinable perceptual process. I'm thinking of cases in which a

³⁹ The correspondence relation between a region of the actual sequence and a region of a possible sequence consists in their both being in the same temporal order relative to other regions of the sequences of which they are a part.

creature has a perceptual process that is radically decomposable. That is, in different situations (were the ball blue instead of red, for example), functionally and physiologically isolated sub-processes kick in. Thus, if, in holding the process fixed, we require that the relevant sequences realize the same super-determinations of determinable natural process-kinds, we'll have to suppose that such a creature is blind.

So what we want – if we want to give an explicit account of holding a process fixed – is a way of picking out natural process-kinds (both for sub-regions of the actual sequence and for the sequence as a whole) at the right level in the determinable-determination hierarchy. It looks like something like the following constitutes the seeing relation:

I see some object *x* by having some visual experience *e* iff there is *any* suitably determinate maximal natural process-kind, *K*, realized by the actual sequence of events that eventuates in *e*, such that for each member, *f*, of some suitable set of features *x* could possess (i) were *x* *f* and were some visual experience, *e*₂, to be the output of a sequence of events that realizes *K*, then *e*₂ would represent something as *F* (and, in so doing, increase the number of things that are represented as *F* by one); and (ii) were *x* not-*f* and were some experience, *e*₃, the output of a sequence of events that realizes *K*, then *e*₃ would not represent *x* as *f*.

What we need, if we are to satisfy the critic's desire for a more explicit account, would be an account of what constitutes a maximal natural process-kind as *suitably* determinate. Sadly, I do not have a proposal in hand that would accomplish this task. But I think we have made some progress.

CONCLUDING REMARKS

It is sometimes supposed that a counterfactual account of perception has no hope. However, the burden of this chapter has been to show how a counterfactual account

might be possible. It must accommodate the challenges posed by the counterexamples to David Lewis' account. It must survive McLaughlin and McGinn's principled objections. It would be nice if its advocate had something useful to say about the issues raised by the generality problem. I have accomplished all of these goals.

The main business of the next chapter is a defense of the Gricean arguments; however, as part of my defense of Grice, I argue that the contents of some visual experiences are fixed by perceptual relations. Thus, if the arguments of Chapter 3 are correct, then the theory presented in this chapter, is part of proto-semantic theory – it tells us what facts constitute our visual experiences as bearing this or that content.

CHAPTER 3 THE POSSIBILITY OF VERIDICAL HALLUCINATION

0.1: INTRODUCTORY COMMENTS

Arguments for the claim that episodes of visual perception (episodes of seeing *x*) are constituted by – and thus analyzable in terms of – experiences that causally depend on *x* in the right way make heavy use of the notion of veridical hallucination. That is, they express the supposition that it is possible for there to be a match between one’s experience and unseen objects. For example, Grice argues for the causal theory of perception by adverting to the possibility of veridical hallucination. (Grice 1961)

In this chapter, I’ll raise and address doubts about whether veridical hallucination is possible. This chapter has two main parts. In his article, “The Particularity of Visual Perception”, Matthew Soteriou argues that admitting the possibility of veridical hallucination has unacceptable consequences. (Soteriou 2000) In the first part of the chapter, I defend the possibility of veridical hallucination against Soteriou’s argument. In the second part of chapter, I address worries about the possibility of veridical hallucination that center on Searle’s account of the content of visual experiences. In neither case will we find reason to reject the possibility of veridical hallucination. Additionally, in the first part of this chapter I find reason to accept an externalist thesis, according to which the perceptual facts make a difference to the contents of visual experiences. More support for this claim is found in the second part of the chapter.

0.2: THE ARGUMENT FROM VERIDICAL HALLUCINATION

Consider the Argument from Veridical Hallucination. This argument is equivalent to an argument Grice gives in “The Causal Theory of Perception”. (Grice 1961) This argument is designed to undermine the claim that it is sufficient for seeing a thing that one’s experience match it perfectly and to prompt the causal theory. Imagine standing in a room with plaid walls. For His own mysterious reasons a powerful being has determined – without regard to the actual pattern of color on the walls – that you be subject to experience of a sort that you would have were you actually seeing the plaid walls of the room. There you are in the room and having the kind of experience you would have were you actually seeing the plaid walls of the room. Do you see the plaid walls of the room? Plausibly not. Why not? Because your experience was not caused in the right way by the plaid walls you do not see them.

Let me formalize this a bit:

- (1) There is a possible situation, **s**, in which:
 - (1.1) you are subject to some experience, **e**, with match-determining character, **M**.
 - (1.2) There is an object, **a**.
 - (1.3) **a** and **M** are such that there is a perfect match between **e** and **a**.
 - (1.4) **a** does not causally explain the truth of 1.3
 - (1.5) In **s**, you do not see **a**.

- So (2) The truth of 1.3 is insufficient to make it true that you see **a**.

- Moreover (3) The truth of 1.5 is probably constitutionally explained by the truth of 1.4.

In *s*, you were veridically hallucinating – that is: (i) since you were not seeing *a*, even though it was as if you were seeing such a thing, you were hallucinating; and (ii) you were *veridically* hallucinating because there was a perfect match between *e* and *a*. Since, the notion of a match is going to be the locus of our worries, let me say a little about matching.

What is it for an experience to match a thing? We should probably understand matching as follows: If how things are and how they visually seem to be are the same, then the visual experience that constitutes things visually seeming such-and-such matches the way things are. This is because how things appear to be is a function of some feature of one's experience – I have called it the *match-determining character* of experience.

It is natural to understand the match determining character of experience in representational terms. Thus, the match-determining character of experience is determined by the contents of one's experience.⁴⁰ If we understand the match-determining character of experience in terms of contents, then for an experience to match an object is for it to correctly – that is, truly – represent that object. A perfect match would consist in representing some and only things that are true of that object. A veridical hallucination consists in being subject to sensory experience that represents some and only things that are true of an object (hence veridicality) and yet not seeing that object (hence hallucination).⁴¹

⁴⁰ There are, of course, difference possible views about what determines the contents of one's experiences. Thus, for example, sense-data theorists like Jackson and Robinson take the contents of one's experience to be determined by patterns of resemblance between sense-data, awareness of which constitutes one's experience, and worldly objects. (Jackson 1977) (Robinson 1994)

⁴¹ When I speak without qualification of an experience being veridical, I mean that it is entirely veridical. That is to say, it is not to any extent non-veridical.

If veridical hallucination is impossible, then the Argument from Veridical Hallucination is unsound. In particular, 1 is false because 1.3 and 1.5 are not co-possible.

PART 1 SOTERIOU'S CHALLENGE

1.1: FROM VERIDICAL HALLUCINATION TO VERIDICAL MISPERCEPTION

In his article, "The Particularity of Visual Perception", Matthew Soteriou argues that admitting the possibility of veridical hallucination has unacceptable consequences. Soteriou argues that if you accept the possibility of veridical hallucination, then you must accept the possibility of veridical misperception. Soteriou argues that since the concept of veridical misperception is incoherent, we should reject the possibility of veridical hallucination. (Soteriou 2000)

What I will do in this part is argue that the intended inference from the possibility of veridical hallucination to the possibility of veridical misperception is invalid. First, I'll present Soteriou's argument that veridical hallucination is impossible. Next, I'll address Soteriou's argument by describing a theory of the content of visual experiences that is consistent with the data Soteriou invokes, but allows for the possibility of veridical hallucination. Finally, I'll show how this response causes trouble for a second Gricean argument – one that supposes that it is possible for an experience to match two objects, one of which is seen and one of which is unseen. If arguments based on the possibility of veridical hallucination succeed, then arguments based on the possibility that an

experience might match two objects, one seen and the other unseen, fail. We will save one Gricean argument, but lose another.

The structure of Soteriou's argument is perspicuously laid out in the following two passages:

- A** If an account of the content of experience allows for the possibility of veridical hallucination, then that account allows that the question of the veridicality of an experience can be settled independently of whether an object is being perceived. (Soteriou 2000:177)
- B** If we allow that the question of the veridicality of a subject's experience can be settled independently of the question of whether there is an object being perceived, we thereby allow that the question of the veridicality of an experience can be settled independently of which particular object is being perceived. And if we allow that the question of the veridicality of an experience can be settled independently of the question of which object is being perceived, we thereby allow for the possibility of veridical misperception. (Soteriou 2000:179)

That veridical misperception is an incoherent notion can be easily seen (provided we make some plausible assumptions). Veridicality is a property that an experience possesses just in case it is accurate – where accuracy can be understood as truth. Thus an experience is veridical if the way it represents things to be is the way they are. To misperceive is to see a thing, but in seeing it, be subject to experience that represents it falsely. Thus a veridical misperception would be an experience that is, impossibly, both veridical and non-veridical.

Since this is impossible, anything that would entail veridical misperception must be rejected. Thus, according to Soteriou's argument, since the possibility of veridical hallucination lies at the head of this chain of entailments, it is ruled out by the incoherence of veridical misperception.

Let's start at the end. Soteriou claims that if veridical hallucination is possible, then so is veridical misperception. Since veridical misperception is clearly impossible, so is veridical hallucination. Actually, this talk of veridical misperception is a bit confusing. In essence, the argument is this: there are cases advocates of the possibility of veridical hallucination must count as veridical, but persuasive intuitive considerations would count as non-veridical.

Since the middle terms of Soteriou's argument (the propositions regarding settling questions about veridicality) can be a bit obscure, I shall begin by constructing an argument from the possibility of veridical hallucination to the possibility of veridical misperception. This argument will tell us what would be needed to get from the possibility of veridical hallucination to the possibility of veridical misperception.

Consider the Case of the Many Blue Dots (call it *MBDI*): there is a white wall and on it is a horizontal series of qualitatively identical blue dots. They are all the same size and color and equidistant. You are positioned not far from it and peering into a binocular device – maybe it looks like the kind of thing ophthalmologists might use to administer a test of visual acuity. It is as if you see four blue dots. As it happens, the device contains special lenses. The lenses can, at the push of a button, be either retracted or dropped into position. The position of lenses makes a difference to which of the many dots on the wall you see. If the lenses are retracted, you see the portion of the series of blue dots that is right in front of you in your normal field of vision. If the lenses are in place, then you see a different set of dots – those a foot to the right. However, whether the lenses are in position or not makes no difference to how things look. The spacing of the dots and the shift produced by the lenses are such that whether the lenses are position

makes no difference to how things look. Thus, if the lenses were retracted, it would also be as if you see four blue dots. Suppose that as you peer at the wall through the device, the lenses are in position – they are shifting light. Is your experience veridical or not?

A strong intuitive case can be made that your experience in this case is non-veridical. Consider a situation differing from MBD1 such that, instead of a series of blue dots, there is only one blue dot. Now, as you look through the device, lenses in place, you see the blue dot, but due to the distorting effect of the lenses the dot appears to be a foot to the right. It seems clear that your experience in this case is non-veridical. The dot appears to be somewhere it isn't. Thus, at least some part of how things appear to be is not how things really are.

But notice that dot-by-dot the same thing is true in the case in which there are many dots on the wall. Considered alone, the relation you bear to each dot is such that were it the one dot on the wall you would misperceive its location. Each one appears to be somewhere it isn't. Each dot that appears before you appears to be a foot to the right of its actual location. Since this is so, it seems we should conclude that as you peer through the device, lenses in place, the experience you have is non-veridical. Your experience in MBD1 is not veridical.

However, according to Soteriou anyone who accepts the possibility of veridical hallucination must think your experience in MBD1 is veridical. What view of the content of experience would make it true that your experience in MBD1 is veridical? Such a view would have to require that, in this case, the veridicality of your experience does not depend on the perceptual facts. What I mean is this: our reason for thinking that your experience was non-veridical depended on thinking that the veridicality of your

experience depended on the states of some dots and not others. The experience misrepresented because the dots you saw were not where they were represented to be.

This means that if a theory is to count your experience in MBD1 as veridical, that theory must require that the veridicality of your experience in MBD1 not depend on the locations of seen dots in MBD1. If MBD1 is veridical, its veridicality would have to depend on the state of unseen objects.

Let's back up a moment and see whether the possibility of veridical hallucination implies such a theory. If it is possible to suffer from a veridical hallucination, then it is possible that an experience be such that it is hallucinatory and yet veridical at the same time. Just as your experience in MBD1 could not depend for its veridicality on the location of seen objects (if it were to be veridical), so the experience you might have a result of hallucination could not depend on the features of seen objects (if it is to be veridical). This is because it is a hallucinatory experience in which nothing is seen. So the content of a veridical hallucination and the content of an experience that could be veridical in MBD1 would both determine veridicality conditions that do not depend on the features of seen objects. They would depend on the states of objects – whether they are seen or not.

If accepting the possibility of veridical hallucination is to commit one to accepting that veridical misperception is possible, then it must do so via some principle that generalizes from the case of veridical hallucination to cover at least cases like MBD1. That is, there must be some principle that allows one to generalize from the claim that the experiences that are parts of veridical hallucinations do not depend for their veridicality

on the states of seen objects to the claim in cases like MBD1 experiences do not depend for their veridicality on the states of seen objects, but rather on unseen objects.

Suppose (as is plausible) that the experience you have in MBD1 could have been part of an episode of veridical hallucination. If its contents would have been the same were it an episode of veridical hallucination, then just as a veridical hallucination has contents in virtue of which its veridicality does not depend on the states of seen objects so would the experience you have in MBD1.

But why suppose that the experience you had in MBD1 and its veridical hallucinatory counterpart share exactly similar contents? A principle that answers this question would get us from accepting the possibility of veridical hallucination to accepting the possibility of veridical misperception.

At this point, it is worth adverting to Martin Davies. According to Davies, “[the notion of] the perceptual content of experience is a phenomenological notion.” (Davies 1997: 314) The root idea is this: the theoretical function of postulating contents for visual experiences is to account for phenomenological similarities and differences among possible visual experiences.

Phenomenally indiscriminable things “give rise to the same perceptual contents”. (McGinn 1989: 66) As Davies puts it, “where there is no phenomenological difference for the subject, there is no difference in content.” (Davies 1997: 314) If two possible experiences are phenomenological twins, then they have the same content. Similarly, presumably, if two possible experiences are not phenomenological twins, they have different contents. Call the conjunction of these two views the *Phenomenological Content Thesis*.

It is plausible that the perceptual facts do not by themselves make a phenomenological difference because you can readily imagine that the experience to which you are subject right now differing with respect to the perceptual facts, but remaining phenomenologically identical. Though I now see this desk, I could have been viewing a qualitatively identical but numerically distinct desk. I could have been the victim of total hallucination. In each case, my experience would be phenomenologically the same. This is because in imagining that this experience has a different perceptual status, I hold fixed the phenomenal character of the experience. This is what makes the imagined scenarios – scenarios in which the perceptual facts are different – scenarios in which this experience exists. Accordingly, by the Phenomenological Thesis, in each case the contents of my experience would be the same. Thus, the perceptual facts do not make a difference to content. Generally, if an experience were different with respect to its perceptual status, then, were the phenomenal character of the experience held fixed, its content would remain the same.⁴²

Therefore, given that the experience you have in MBD1 could have been a phenomenologically identical veridical hallucination, it must share exactly similar contents with that veridical hallucination. Given that the veridical hallucination didn't depend for its veridicality on the status of seen objects, the same will be true of your experience in MBD1. Thus the only obstacle to your experience being veridical in

⁴² Actually, this is a little weaker than would be implied by the Phenomenological Thesis. But that is both acceptable and a good thing. It is acceptable because it doesn't need to be any stronger to warrant the generalization from cases of veridical hallucination to MBD1. It is a good thing because rampant externalism about what features a mental state represents would falsify the Phenomenological Content Thesis, if it were interpreted to apply to possible twins from distant possible worlds (e.g. experiences on Earth and Twin Earth). Luckily, this claim has a plausibility independently of the Phenomenological Thesis.

MBD1 is removed and, if we accept the possibility of veridical hallucination, we must accept the veridicality of your experience in MBD1.

Thus Soteriou places the advocate of the possibility of veridical hallucination in a cognitive bind – the advocate must live with inconsistency, reject the possibility of veridical hallucination, reject the extremely persuasive intuitive verdict that your experience in MBD1 is non-veridical, or reject the reasoning that leads from the possibility of veridical hallucination to the possibility of veridical misperception. I propose the last path.

1.2: ASSESSING SOTERIOU'S ARGUMENT

The argument contained in passages **A** and **B** is invalid. Soteriou claims that the possibility of veridical hallucination entails that (1) “the question of the veridicality of an experience can be settled independently of whether an object is being perceived.” (Soteriou 2000:177) The claim that the question of the veridicality of an experience can be settled independently of whether an object is being perceived is susceptible to at least two readings. On one it makes only a weak claim – that there *some* experiences are such that their veridicality at a situation does not depend on whether an object is being perceived in that situation. On another it makes a much stronger claim – that *every* experience is such that its veridicality at a situation does not depend on whether an object is being perceived in that situation. Charity would recommend the weak reading of (1) since it is implied by the possibility of veridical hallucination. But unless we give (1) a strong reading it won't entail the possibility of veridical misperception. Sadly, however, the strong reading is not entailed by the possibility of veridical hallucination.

B contains the claim that (2) “the question of the veridicality of an experience can be settled independently of which particular object is being perceived.” (Soteriou 2000: 179) Like (1), (2) admits a strong and weak reading. On the weak reading (2) requires that *some* experiences are such that their veridicality at a situation does not depend on which object is being perceived in that situation. On the strong reading (2) requires that there *every* experience is such that its veridicality at a situation does not depend on which object is being perceived in that situation.

B contains the claim that (2) follows from (1). It is plausible that the weak reading of (2) follows from both readings of (1). However, the strong reading of (2) does not follow from the weak reading of (1). It could be true that some experiences are such that their veridicality is independent of whether something is being perceived but false that every experience is such that its veridicality is independent of which objects are perceived in the perceptual episode of which it is a part.

We need a strong reading of (2) in order that (2) might entail the possibility of veridical misperception. If *every* experience is such that its veridicality at a situation does not depend on which object is being perceived in that situation, then in cases like MBD1 we must count the experience you have as veridical. But persuasive intuitive considerations support the claim that the experience you have in MBD1 is non-veridical. Thus, a strong reading of (2) would place us in a cognitive bind. The weak reading of (2) won't do the trick, since if only some experiences are such that their veridicality is independent of which object is perceived, then there is a possibility that the experience you have in MBD1 is not among those experiences the veridicality of which is independent of which objects are perceived.

Unfortunately for Soteriou, while it is plausible that the strong reading of (2) follows from the strong reading of (1), it is not plausible that the strong reading of (1) follows from the possibility of veridical hallucination. All the possibility of veridical hallucination requires is that some experiences – namely those that are parts of episodes of veridical hallucination – are such that their veridicality is independent of whether anything is being perceived.

Accordingly, Soteriou has failed to establish that the possibility of veridical hallucination entails the possibility of veridical misperception. However, as we saw before in my analysis of MBD1, if the accepting possibility of veridical hallucination is to commit you to the possibility of veridical misperception, you must also accept some kind of generalization principle – a principle that gets you from *some* to *all* (or at least from some cases to cases like MBD1).

Whatever generalization principles one chooses will come under suspicion from its implication (together with the possibility of veridical hallucination) that your experience in MBD1 is veridical. The dialectical situation seems to me to be this: veridical hallucination seems possible, persuasive considerations establish that in MBD1 your experience is non-veridical, therefore we have reason to reject whatever principles get us an entailment that your experience in MBD1 is veridical. Thus, it seems that we ought to accept the possibility of veridical hallucination and reject the generalization principles, whatever they are.

It is surprising that Soteriou supposes that the possibility of veridical hallucination must be given up, since all the pieces are in place in his article for a view that would accept the possibility of veridical hallucination, while denying the veridicality of your

experience in MBD1. The view I have in mind is one that denies the claim that if an experience were different with respect to its perceptual status, then, were the phenomenal character of the experience held fixed, its content would remain the same, while accounting for the evident fact that the veridicality of your experience in MBD1 depends on the perceptual facts. I'll present this view by showing how it diagnoses a failed argument that your experience in MBD1 is veridical.

A failed argument: Consider looking through the device when the lenses are retracted. In this case, call it *MBD2*, how things are and how things look are the same (if they ever are). There is a way things look, and it is the same as the way things are. There appears to be a white wall with some blue dots arrayed on it thus-and-so and there is a white wall with some blue dots arrayed on it thus-and-so. The experience you would have in this case is therefore veridical. Since things look the same whether the lenses are in place or not, the same facts obtain in MBD1. So, since there appears to be a white wall with blue dots arrayed on it thus-and-so and there is, the experience you have in MBD1 must also be veridical.

This argument supposed that the contents of one's experience are determined by the things-looking-*X*-state one is in being subject to that experience. It is in virtue of being in such-and-such experiential state that things look as they do. Moreover, it is plausible that the facts about how things look mirror the contents of the experiences involved in things looking the way they do. Think of the mirroring relation here as follows: an experience possesses only those contents necessary to capture how things look in being subject to it.

The differences that exist between MBD1 and MBD2 are not sufficient to make a difference to how things look. Whether the lenses are in place or not makes no difference to how things look. If the contents of experiences mirror how things look, then the position of the lenses do not matter to whether your experiences in MBD1 or MBD2 are veridical. Accordingly, if your experience in MBD2 is veridical, so is your experience in MBD1.

I suspect that part of what explains the persuasive force of the argument that one's experience in MBD1 is veridical is a kind of sloppiness in our notion of things looking thus-and-so. In particular, I said that whether the lenses are in position or not would make no difference to how things look. This is true.

You might think: if, in two situations, things look the same to someone, then the contents of his visual experience in those two situations must be the same. However, this thought expresses an invalid inference. The perceptual facts (which objects, if any, someone is seeing) might make no difference to *how* things look, but they might make a difference to *which* things look thus-and-so.

If I ask you how things look, my question does not necessarily target the identities of seen objects, but rather the features they appear to have. It is an unspecific question. You wouldn't answer an inquiry into how things look with an utterance of "Fred". Instead, you might say "blue", "blurry", "as if Sally finally had it with Charles", "as if there were a ghost standing behind that tree", "as if something purple and hairy were located inches from my face", etc. The question of how things look can be answered with radically varying degrees of specificity about which objects one bears perceptual relations to.

The intuitive considerations in favor of the conclusion that in MBD1 your experience is non-veridical rest on the idea that the perceptual facts make a difference to which things look-thus-and-so. Your experience in MBD1 is non-veridical because the objects you see (the objects providing answers to the question “which objects do you see?”) are not as they appear to be. Had you been aware of different objects (as in MBD2) the veridicality of your experience would have depended on those objects instead. Thus, as Soteriou puts it, “We need to determine which particular objects in the subject’s environment are being perceived if we are to determine whether the subject’s environment really is as it seems to be”. (Soteriou 2000: 180) A proper understanding of MBD1 and MBD2 indicates that at least sometimes it matters to what the contents of one’s visual experience are not just how things look, but *which* things look *how*, as it were.⁴³

We need not reject the claim that there is a mirroring relation between looks-states and the contents of experience. These considerations allow that there might be such a relation. Instead, they suggest that we should reject the claim that the relevant looks-states are states of things-looking-X. Rather the mirroring relation would obtain between the contents of experiences and states in which particular things look X.

Since the perceptual facts make a difference to the contents your experience it is false that if an experience were different with respect to its perceptual status, then, were the phenomenal character of the experience held fixed, its content would remain the same. The perceptual facts play a role in fixing the contents of visual experiences. The contents of the experiences you have in MBD1 and MBD2 differ because the experiences

⁴³ The inference to this sentence from the previous one is a substantial one and depends on rejecting Searle’s theory of the contents of visual experiences. However, I’d like to table a discussion of Searle for the time being. If you cannot wait, please consult Part 2 of this chapter for much more.

possess different perceptual statuses. In MBD1, the experience is part of an episode of perceiving one set of blue dots. In MBD2, the experience is part of an episode of perceiving another set of blue dots. Accordingly, the veridicality of the two experiences depends on the states of different objects.

Properly understood to be non-veridical, cases like MBD1 warrant the conclusion that when you are seeing something, the content of your experience requires that the seen thing possess such-and-such features if your experience is to be veridical. Consequently, they warrant the conclusion visual experiences have different contents in different perceptual circumstances.⁴⁴ In particular, they warrant the conclusion that two different perceptual experiences might differ in their contents, only because they are embedded in different perceptual relations.

But what should we think about hallucinatory experiences? Once we have allowed that the contents of an experience might vary with changes in perceptual status we open the possibility that hallucinatory experiences might typically have contents that differ from the contents of their perceptual counterparts. All we need to allow for the possibility of veridical hallucination is that a hallucinatory experience can depend for its veridicality on the states of unseen objects. This does not tell us very much about the behavior of contents of hallucinatory experiences.

One hypothesis that is consistent with current constraints is this: hallucinatory experiences have contents that satisfy the Variety Condition. An experience satisfies the Variety Condition if and only it can be made true, its content held fixed, by a variety of possible objects. For example, suppose that your experience, as you look at Fred, satisfies the Variety Condition. It follows that while the state of Fred might make your

⁴⁴ Again, the implication here depends on the falsehood of Searle's view.

experience veridical, the state of Fred's qualitative twin could also have made that experience veridical. Such contents prescind from the synchronic identities.⁴⁵

On this account, hallucinatory experiences possess only the kind of contents McGinn and Davies think all visual experiences possess. In doing so they affirm what is sometimes called the *Generality Thesis*. McGinn claims that "the content of experience is not to be specified by using any terms that refer to the object of experience." (McGinn 1982: 51) This claim is sometimes explicated by invoking the idea of quantification: "...we can take the content of perceptual experience to be existentially quantified content. A visual experience may present the world as containing *an* object of a certain size and shape, in a certain direction, at a certain distance from the subject" (Davies 1992: 26)

His acceptance of the claim that every experience satisfies the Variety Condition explains McGinn's talk of existential quantification. Sentences containing existential quantification bear on particulars in a way that is analogous to the way an experience satisfying the Variety Condition bears on its objects. Consider, for example, the sentence:

(S1) The red-haired girl mowed the lawn.

On Russell's analysis, S1 has the same meaning as:

(S2) There is exactly one relevant red-haired girl and she mowed the lawn.

Now given that it is contingent which among all the possible things is a relevant red-headed girl, S1 and S2 could be made true by ever so many things. Even if you hold

⁴⁵ Experience can prescind from synchronic identities while not prescinding from diachronic identities.

fixed the contents of sentences like S1 and S2, their being about one thing rather than another is contingent.⁴⁶ Suppose that in the actual world there is exactly one relevant red-headed girl, Sally, and she mowed the lawn. Suppose that in another world there is exactly one relevant red-headed girl, Frieda, and she mowed the lawn. What I would have said by an actual utterance of S2 would be true at the actual world and at the possible world in which Frieda mowed the lawn. Now we can see the point of claiming that visual experiences have quantificational contents. Just as S1 and S2 can be made true by Frieda, they can also be made true by Sally. That this is so is explained by the semantic function performed by existential quantification. Likewise, since the contents of visual experiences satisfying the Variety Condition can be made true by a variety of possible objects it is not surprising that people use talk of existential quantification to express the Generality Thesis.

So while I doubt that all the contents of every experience satisfy the Variety Condition – in particular, I doubt that perceptual experiences satisfy the Variety Condition – it is possible that hallucinatory experiences possess contents that do. Thus it is possible that hallucinatory experiences have contents in virtue of which veridical hallucination is possible. Your experience as of a plaid wall, when hallucinatory, might be made true just in virtue of the fact that there is a plaid wall where it seems there is a plaid wall.

⁴⁶ Whether the same is true of sentences containing rigidified descriptions in their quantificational phrases would require looking into the difference between counterpart theory and theories according to which there is a primitive relation of transworld identity.

1.3: THE DEFENSE OF VERIDICAL HALLUCINATION AND THE ARGUMENT FROM COMPETING OBJECTS

My solution to the problem posed by Soteriou's Argument raises trouble for the second chief Gricean argument, the Argument from Competing Objects. Like the Argument from Veridical Hallucination, the Argument from Competing Objects is both destructive and constructive. Grice invokes cases in which one's experience matches two objects – only one of which is seen. These cases are supposed to establish that experience-object matches are insufficient for seeing a thing and warrant the hypothesis that causal dependence is a partial constituent of the seeing relation.

The *prima facie* problem is this: what allowed for the possibility of veridical hallucination was the possibility that hallucinatory experiences possess contents that can be made true by a variety of possible objects. Cases like MBD1 suggest that the same is not true of perceptual experiences. That is, cases like MBD1 suggest that perceptual experiences are not such that a variety of possible objects could make them true. The worry here is that the Argument from Competing Objects is unsound. It supposes that something impossible is possible.

As it will turn out, what attitude we should take about the Argument from Competing Objects depends on the relationship between the contents of hallucinatory experiences and the contents of perceptual experiences.

Let's look in some detail at the Argument from Competing Objects. I'll start with Grice's statement of it:

...it might be that it looked to me as if there were a certain sort of pillar in a certain direction at a certain distance, and there might actually be such a pillar in that place; but if, unknown to me, there were a mirror interposed between me and

the pillar, which reflected a numerically different though similar pillar, it would certainly be incorrect to say that I saw the first pillar, and correct to say that I saw the second. (Grice 1961:)

Formalizing a bit (and changing from 'I' to 'you') we get:

- (1) There is a possible situation, **s**, in which:
 - (1.1) you are subject to some experience, **e**, with match-determining character, **M**.
 - (1.2) There are two objects, **a** and **b**.
 - (1.3) **a** and **M** are such that there is a perfect match between **e** and **a**.
 - (1.4) **a** does not causally explain the truth of 1.3.
 - (1.5) **b** and **M** are such that there is a perfect match between **e** and **b**.
 - (1.6) **b** does causally explain the truth of 1.5.
 - (1.7) In **s**, you do not see **a**, but you do see **b**.
- So (2) The truth of 1.3 is insufficient to make it true that you see **a**.
- Moreover (3) The truth of 1.7 is probably constitutionally explained by the truth of 1.4 and 1.6.

Thus, Grice argues that it is insufficient for seeing a thing that your experience match it. What you need, in order that you might see a thing, is that your experience causally depend on it. Thus we get the causal theory of perception.

The worry here is that 1.5 and 1.7 are not genuinely co-possible. If so, 1 is false and the Argument from Competing Objects is unsound. Remember, we're understanding matching in terms of veridicality and veridicality in terms of truth. Thus the Argument from Competing Objects depends on the supposition that it is possible for a perceptual experience to be made true by two distinct objects, one of which is seen and one of which is unseen. The denial that 1.3 and 1.5 are co-possible is the denial that it is possible for a perceptual experience to be made true by two distinct objects, one of which is seen and one of which is unseen.

What grounds are there for thinking that a perceptual experience could not be made true by two distinct objects, one of which is seen and one of which is unseen? The argument here is simple. Consideration of MBD1 suggests that perceptual experiences have contents that do not satisfy the Variety Condition. Moreover, MBD1 suggests that whether a perceptual experience is veridical depends on what features the objects of the perceptual episode of which it is a part possess. Now, suppose that perceptual experiences have *only* contents that do not satisfy the Variety Condition and suppose that these contents do not bear on unseen objects.⁴⁷ If so, it will be impossible for an experience to be made true by two distinct objects, one of which is seen and one of which is unseen.

It might be helpful to consider a linguistic analogy here. Consider the following sentence:

(S4) Claudius is status-conscious.

Whether S4 is true depends on the state of Claudius. Whether it is true at a world depends on Claudius' state at that world. He is the truth-maker for S4. Now what about Fred? Suppose Fred is also status-conscious. Does S4 match Fred? Holding the content of S4 fixed, Fred's status-consciousness has no direct bearing on the truth of S4. Since S4's content does not satisfy the linguistic analog to the Variety Condition, it cannot be made true by Fred. The worry about the Argument from Competing Objects stems from the thought that the contents of perceptual experiences are like the content of S4. They cannot be made true by a variety of possible objects.

⁴⁷ Ignoring, for the time being, the bearing a visual experience might have on the unperceived perceiving subject.

The problem with the Argument from Competing Objects is that if **b** is seen and **a** is not, then there cannot be a match between **a** and **e**. Giving this a representational gloss, **a** cannot make **e** true. 1.3 is false if 1.7 is true. We therefore have reason to think the Argument from Competing Objects is unsound.

1.4: SHIFTING CONTENTS: REPLACEMENT AND ADDITION

Before I evaluate the objection to the Argument from Competing Objects, I'd like to return for a moment to my defense of the Argument from Veridical Hallucination. The move I made there was to raise the possibility that, while the contents of perceptual experiences do not satisfy the Variety Condition, the contents of at least some hallucinatory experiences do. Such hallucinatory experiences would then be apt to match a variety of possible objects (thereby allowing for the possibility of Veridical Hallucination). The contents of an experience shift as its perceptual status shifts.

But I have yet to say much about what this shift consists in. In particular, the shift could consist either in replacement or in addition. The shift in contents consists in an *addition* if all experiences possess contents that satisfy the variety condition (and thus, given these contents could be veridical hallucinations), but perceptual experiences also possess some contents that do not satisfy the Variety Condition. The shift in contents consists in *replacement* if only hallucinatory experiences possess contents satisfying the Variety Condition, while their perceptual counterparts possess contents that do not satisfy the Variety Condition. Whether shifting is effected by replacement or addition, there is trouble for the Argument from Competing Objects.

If the shift consists in replacement, then the problem with the Argument from Competing Objects is acute. If the shift consists in replacement, then no perceptual experience that is part of an episode of seeing one thing could be veridical at all with respect to another, unseen thing. That is, no perceptual experience of x could be made even in part true by y – where y is in fact unseen.

If the shift consists in addition, then no perceptual experience that is part of an episode of seeing one thing could be veridical without qualification with respect to another, unseen thing. That is, no perceptual experience of x could be made entirely true by y – where y is in fact unseen. Some of the contents of the perceptual experience of x could be made true by y – those contents satisfying the Variety Condition. However, other contents cannot be made true by y – those contents that, in virtue of the perceptual facts, can only be made true by x . Thus, if the shift consists in replacement, then not every content of a perceptual experience could be made true by unseen objects. Of course, for all this, the view that shifting consists in addition, allows that some of the contents of a perceptual visual experience might be made true by unseen objects.

Now, given the argument of the previous two paragraphs, it ought to be surprising that we can imagine what is, apparently, impossible – that a perceptual experience might match two objects equally. If this is not really possible, then what are we imagining?

One might hold out hope that there is room to maneuver if shifting is addition. We can, I think, conceive of an experience without regard to whether or which object is perceived in the episode of which it is a part. When we do so, we can have some grip on its contents – we can conceive of it as possessing such contents as satisfy the Variety Condition. That we can do this suggests that the view that shifting is addition is correct.

There is some subset of the contents of a perceptual experience that can be held fixed – and assessed for veridicality relative to multiple objects – even as we suspend judgment about which perceptual facts obtain. Thus, while an experience might not be such that all its contents are made true by a variety of possible objects, perhaps when we contemplate the possibility envisioned in the Argument from Competing Objects we pick out the contents of a visual experience that can be made true by a variety of possible objects – namely, those contents that are shared with the hallucinatory counterparts of these perceptual experiences.

Perhaps, then, we can imagine an experience that matches two objects equally, while also imagining that the experience is a perception of one but not the other (though remaining in temporary suspense about which object is perceived). Presumably this imaginative project determines a possibility of some sort – the possibility that an experience might possess some contents in virtue of which two objects match it, only one of which is seen. However, this possibility is of no use to Grice. It doesn't establish that matching is insufficient for seeing, at best it establishes that a partial match is insufficient for seeing. If the case for the insufficiency of matching accounts of seeing rested entirely on the Argument from Competing Object, then it would not succeed.

1.5: THE POSSIBILITY THAT INVOKING CAUSAL DEPENDENCE IS GRATUITOUS

The response I gave to the Soteriou's argument that veridical hallucination is impossible, might lead one to think that invoking causal dependence – or some close cousin of it – is gratuitous. If our theoretical goal in giving an account of the perceiving relation is to discover some feature that distinguishes perceptions from hallucinations or

perceptions of one thing from perceptions of another, then, given how I suggest the perceptual facts relate to facts about the contents of experience, we can find success immediately. In honor of Boris Kment, who first suggested such a view to me in conversation, I'll call this *Kment's Suggestion*.⁴⁸

If there is some kind of content that all and only perceptual experiences – that is, experiences that are parts of episodes of seeing – possess, then it is possible to give an extensionally adequate account of the distinction between seeing and hallucinating. Episodes of seeing could be distinguished from episodes of hallucination by their possession of this kind of content.

While hallucinatory experiences have only contents that satisfy the Variety Condition, perceptual experiences have contents that do not satisfy the Variety Condition. As it stands, such an account would distinguish cases of seeing from cases of hallucination, though it wouldn't distinguish cases of seeing one thing from cases of seeing another thing. It might, however, be readily extended to distinguish such cases. Seeing *x* is distinguishable from seeing a distinct thing, *y*, insofar as *x* is the relevant truth-maker for the experience that is part of seeing *x*, while *y* is the relevant truth-maker for the experience that is part of seeing *y*. Episodes of seeing *x* could be distinguished from episodes of seeing *y* by the contents of the experiences involved in each.

There are two problems for this line of thought. First, it is not clear that hallucinatory experiences must possess only contents satisfying the Variety Condition. Hallucinatory experiences of the sort invoked in cases of veridical hallucination are not hallucinations of particular objects (even if, to their victims, they seem to be perceptions

⁴⁸ Of course, he gave no indication of believing that Kment's Suggestion represented his actual view of the matter.

of particular objects). An external observer (someone who is not subject to the hallucination) would not give a *de re* report of the experience – he would not say “Sally [or something the external observer takes to be a denizen of the actual world] is such that Fred [the victim of the hallucination] experiences her to be brown-haired”. Such hallucinatory experiences have objects only in the sense in which the fountain of youth was the object of Ponce De Leon’s quest.

However, if I had to guess, actual hallucinations frequently have particular things as their objects. Someone can, for example, hallucinate his own mother.⁴⁹ In describing this person’s predicament, one might say “Fred’s mom is such that Fred experiences her to be brown-haired.” It is plausible that such hallucinations involve experiences that depend for their veridicality on the states of the existing hallucinated object. Some support for this comes from the possibility that someone might mis-hallucinate his own mother – in being subject to experience that represents her as brown-haired when really she is red-headed. It is plausible that this is what is going on when someone reports a dream by saying things like “I dreamt my mother was there... only it wasn’t my mother (because the woman I dreamed about had brown hair).” That someone might mis-hallucinate his mother suggests that it is possible to have hallucinatory experiences that depend for their veridicality on the state of particular things.⁵⁰ This line of thought

⁴⁹ Mark Johnston adverts to this fact, in a different dialectical context, in his article, “The Obscure Object of Hallucination”. (Johnston 2004: 129)

⁵⁰ A philosopher might hold that it is not, strictly speaking, the content of the hallucinatory experience of one’s mother that entails that it is one’s mother that is the object of a hallucination. Perhaps, for example, all that is given in the content of the experience is that there is some woman who is such-and-such and the subject takes that woman to be his mother. What constitutes the episode as a hallucination of the mother is an intellectual act – the act of judging that the thing that is such-and-such (as given in experience) is the mother. In other words, the of-ness of a hallucination is not the of-ness of the contents of the experience that is part of it. On this line of thought, the phenomenon I’ve called ‘mis-hallucination’ would have to be understood as a hybrid state. That is it consists of a hallucinatory experience the content of which, like the contents of veridical hallucinatory experiences, satisfies the Variety Condition, together with a judgment

supports the claim that we cannot generally distinguish hallucinatory experiences from non-hallucinatory experiences by their contents. It is not a necessary condition on being a hallucinatory experience that it possess only contents satisfying the Variety Condition.

Suppose we set aside this first doubt about Kment's Suggestion. Even then, there is a problem. The goal is not just to pick some feature that all and only perceptual experiences of some particular possess, it is also to discover what constitutes an episode as an episode of perceiving some particular. If the arguments of Chapter 2 are correct, then among the features all and only episodes of perception possess is the feature of involving an experience that counterfactually depends on states of objects. If the line of resistance to Soteriou's argument is correct, then among the features all and only episodes of perceiving *x* possess, is a content that depends for its veridicality on *x*. At the very least, if we follow Kment's Suggestion, we'll have an incomplete theory of perception.

However, I'd like to argue for a stronger conclusion: the fact that experiences possess the kinds of contents I've described in this section, is a mere consequence of the perceptual facts, rather than a constituting determinant of the perceptual facts. The problem with supposing, as Kment's Suggestion suggests, that facts about the contents of experience are constitutive determinants of the perceptual facts is that it rules out certain kinds of permissible explanation. It is an interesting question: why does this experience have this kind of content – why does its veridicality depend on Elizabeth rather than

that the thing apparently perceived is the mother. The misrepresentation of the mother occurs not in experience but in judgment – in judging (based on the experience) that the mother is brown-haired.

The possibility of *radical* mis-hallucination lowers the odds that the view described in this footnote is correct. One might mis-hallucinate one's mother and in doing so misrepresent her almost entirely. One might have a dream in which one's mother was a fragile vase that required transport through crowded streets. Now, according to the hybrid account of mis-hallucination, this would have to be a case in which one judged, based on one's experience, that the vase was one's mother. However, since this judgment would be deeply crazy (if the experience did not already purport to reveal your mother), it is preferable to reject the hybrid view and claim that it was the content of your experience that was crazy and not your faculty of judgment.

Elaine? It is permissible to answer the question by adverting to the perceptual facts. But if the fact that the veridicality of an experience depends on Elizabeth rather than Elaine constituted an experience as part of an episode of perceiving Elizabeth, then this explanation would be impermissible. It would purport to explain something by covertly invoking that very fact.

1.6: CONCLUDING REMARKS

In defending the possibility of veridical hallucination from Soteriou's argument I suggested that the perceptual facts play a role in determining the veridicality conditions of experiences. That an experience is hallucinatory allows that it might have contents that satisfy the Variety Condition. That an experience is part of an episode of seeing something makes it true that its veridicality depends on the states of the seen thing – rather than the states of unseen things. I have argued that the perceptual facts play a proto-semantic role – that is, they are among the constitutive determinants of the contents of visual experiences. This allowed us to resist Soteriou's push from accepting the possibility of veridical hallucination to accepting the possibility of veridical misperception.

It has the consequence that whatever principles would get one from accepting the possibility of veridical hallucination to accepting the possibility of misperception must be rejected. Interestingly, this means that insofar as we accept the view I've presented in Part 1 we have reason to reject even fairly weak versions of the thought that the content of visual experience is a phenomenological matter. For example, we have reason to reject the idea that if an experience were different with respect to its perceptual status,

then, were the phenomenal character of the experience held fixed, its content would remain the same. Thus, contrary to what McGinn and Davies hold, the contents of visual experiences are not simply postulates in the attempt to account for the phenomenology of visual experience. The idea of a visual content serves other functions.

I have suggested a strategy for accounting for intuitions about the veridicality of your experience in MBD1, while allowing for the possibility of veridical hallucination. This strategy gives a partial account of the proto-semantics of visual experiences. That is, it gives us some information about what constitutes an experience as possessing contents of a certain sort. This is an externalist theory. Externalists about some feature, F, characteristically claim that whether a creature possesses F depends on facts about how the creature is related to its environment. Your experience in MBD1 depends for its veridicality on the states of the four objects to which you bear the seeing relation. It depends for its veridicality on the states of those objects because the content it possesses depends on the perceptual facts. Since whether you bear the seeing relation to an object depends on your relation to the environment external to you, it follows that externalism about the contents of visual experience is true. The contents of your experience depend on your relation to your environment.

PART 2

SEARLE AND VERIDICAL HALLUCINATION: ARE SEEING RELATIONS REPRESENTED IN THE CONTENTS OF VISUAL EXPERIENCES?

In Part 1 I suggested a partial account of the proto-semantics of visual experiences according to which the veridicality of an experience can depend on the perceptual facts

because the contents of an experience depend on the perceptual facts. Searle would reject this claim and the inferences that support it. On his view, the veridicality of an experience depends on the perceptual facts because the contents of experiences represent the obtaining of those perceptual facts. This means that instead of giving a proto-semantic account of the data, Searle gives a semantic account of the data. Accordingly, on Searle's view, veridical hallucination is an impossibility. Thus, Searle is committed to thinking that the Argument from Veridical Hallucination is unsound. The goal of this section is to present and evaluate Searle's challenge to the possibility of veridical hallucination.

2.1: SEARLE'S ACCOUNT

Searle accepts a version of the Generality Thesis that implies (correctly) that your experience in MBD1 is non-veridical, but rules out the possibility of veridical hallucination. Additionally, Searle's view has the attractive feature of allowing us to give into the temptation to suppose that if an experience were different with respect to its perceptual status, then, were the phenomenal character of the experience held fixed, its content would remain the same.

Searle represents the content of the experience you might have as you look at a yellow station wagon as follows:

Vis Exp (There is a yellow station wagon there and the fact that there is a yellow station wagon there is causing this visual experience). (Searle 1983: 48)

Thus, on Searle's view, visual experiences represent not only distal conditions (how things are with the environment), but also perception-grounding relations between

themselves and these distal conditions. Contrast Searle with the externalist, who might represent the same experience as follows:

Vis Exp (That, there, is a yellow station wagon).

On this externalist picture, experience represents only what might be called ‘distal conditions’. In answering the question of whether the requisite distal conditions obtain we need only ask questions like the following: Is the car yellow? Is it a car? Is it shaped in such-and-such a way? Does it relate to its surroundings in such-and-such a way? Is it such-and-such a distance from you? Etc.

The best way to understand Searle’s theory is to see what analysis of MBD1 it would provide. On a theory of visual content according to which experience represents perception-grounding conditions, the experience you have is non-veridical because the perception-grounding conditions it represents as being fulfilled aren’t fulfilled.

Such content would probably work something like this: Experience represents, first of all, an environment populated with objects at various locations around the viewer. Second, picking out an object by location, it represents itself as causally depending on that object. There appearing to be an object at that location is represented as being causally dependent (in the right way) on that object. Finally, it represents these objects as possessing some sensible features. Thus, visual experiences represent themselves as being caused in the right way by objects located in the environment.

Your experience in MBD1 requires two things in order to be veridical. First, it requires that there be blue dots at some locations – call them L1, L2, L3, and L4. Second, it requires that the appearance of there being a blue dot at L1 causally depend on

a blue dot at L1, the appearance of there being a blue dot at L2 causally depend on the blue dot at L2, and so on. Thus, experiences represent both distal conditions and relations between themselves and those distal conditions.

Searle's account, like the externalist account (according to which the contents of an experience can depend on the perceptual facts) rules out the possibility of veridical misperception. It correctly predicts that your experience in MBD1 is non-veridical – and so is consistent with the judgments of common sense. Your experience turns out to be non-veridical in MBD1 because it requires – for example – that some blue dot at L2 be causally responsible (in the right way) for its own requirement that there be a blue dot at L2. However, since the distorting lenses are in place, the dot at L2 is not responsible for the appearance of there being a dot at L2. Rather, some other dot – e.g. the one at L1 – is causally responsible for the appearance. Your experience is non-veridical because it misrepresents the relationship between your experience and the dots.

How does Searle's account rule out the possibility of veridical hallucination? First, Searle's account is supposed to be entirely general, that is every visual experience is such that it requires for its veridicality that such-and-such distal facts obtain and that perception-grounding relations connect the obtaining of those facts and the appearance that those distal facts obtain. This means that every hallucination misrepresents the facts. Since an episode of hallucination is an episode in which the perception-grounding relations do not connect one's experience and the facts, every hallucinatory experience must be non-veridical. Consequently, any argument that depends on the possibility of veridicality of veridical hallucination must be unsound.

Searle's account also has the consequence that the Argument from Competing Objects is unsound. If it is to be sound, the Argument from Competing Objects requires that an unseen distal object could be a truth-maker for an experience. However, so long as the obtaining of perception-grounding relations to distal objects is among the veridicality conditions of an experience, no unseen distal object could make an experience veridical. An unseen object is such that the perception-grounding relations do not obtain between it and an experience. Thus, if (as Searle requires) a veridical experience bears the perception-grounding relations to its objects, then no experience can be made veridical by a distal object which is also unseen. Consequently, Searle cannot accept the Argument from Competing Objects.

2.2: ASSESSING SEARLE

In this context, whether it is reasonable to accept the possibility of veridical hallucination (and whether, therefore, the possibility of veridical hallucination is to function as a constraint on our account of the contents of visual experiences) depends on the comparative acceptability of my externalist proposal and Searle's account of the content of visual experiences. What reason is there to prefer Searle's account? Is there reason to dis-prefer it?

Searle's main explicit reasons for accepting his theory are its obviousness (once entertained), its utility in solving what Searle takes to be outstanding philosophical problems, and its roots in the Gricean arguments. (Searle 1983: 48-50) Appeals to obviousness are of limited epistemic force and, at any rate, it is not clear that Searle's

account is obvious – given the controversy it has sparked.⁵¹ Let us, then, evaluate Searle’s account in light of more substantial considerations.

It is a bit surprising that Searle references the Gricean arguments, since we have just seen that if Searle’s account is correct, then neither the Argument from Veridical Hallucination nor the Argument from Competing Objects can be sound. Thus, an *ad hominem* case can be made against Searle. His own view is self-undermining in the sense that its acceptance is inconsistent with accepting some of the arguments he cites as its sources. Since Searle gives an explicit presentation of neither the Argument from Veridical Hallucination nor the Argument from Competing Objects, it is possible that his citation of them is an oversight on his part. At any rate, while it is arguable that Searle’s citation of the Gricean arguments is inconsistent with his account (and so irrational for him), it is possible (for all we’ve seen so far) that his account is all-things-considered the best account to take nonetheless.

Searle claims that his account of the contents of visual experiences may be justified by its utility in solving outstanding philosophical problems. Evaluating the details of the use to which Searle puts his account would require going too far afield. Luckily, the primary ‘outstanding philosophical problem’ Searle thinks can be solved by the use of his account of visual content is the challenge posed to internalist accounts of mental contents by advocates of content externalism. (Searle 1983: Ch. 8) Thus we need not evaluate the details of the use to which he puts his account of visual experience, if we can find reasons that undermine either his internalism or the motivations for his internalism.

⁵¹ See the essays in (Lepore & Van Gulick 1991) for a selection of the skepticism Searle’s account has provoked.

In the context of a discussion of externalism about belief, Searle indicates the central role internalism is playing in his thought:

...discussions like this can tend to degenerate into a kind of fussy scholasticism which conceals the basic 'metaphysical' assumptions at issue... Some form of internalism must be right because there isn't anything else to do the job. The brain is all we have for the purpose of representing the world to ourselves and everything we can use must be inside the brain. Each of our beliefs must be possible for a being who is a brain in a vat because each of us is precisely a brain in a vat; the vat is a skull and the 'messages' coming in are coming in by way of impacts on the nervous system. (Searle 1983: 229-30)

It is clear from context that Searle intends this claim to generalize over all of our content-bearing mental states. In particular, it is clear from context that Searle means that for any of our content-bearing mental states, it is possible to be in that very state – with the same content – were we brains in vats. Thus, for example, Searle intends that each of our desires must be possible for a brain in a vat. If each of my desires are possible for a brain in a vat, then, were I hallucinating, rather than seeing, I should be subject to the same desires. Thus, on Searle's line of thought, it counts in favor of his theory that according to it mental states are such that their contents would remain the same even if tokened in a hallucinating subject.

Indeed, the possibility of thinking thoughts with a given content – whether one is a victim of hallucination or not – is among the chief motivations of Searle's descriptivism. On Searle's account, reference to particulars (other than ourselves and our own mental states) is secured by a particular's satisfaction of by conditions the obtaining of which is part of the content of our mental states. (Searle 1983: 17) This has the advantage of allowing us to make sense of cases of reference failure without either postulating Meinongian entities to serve as the subjects of our thoughts or concluding that thoughts in which there is reference failure differ in their contents from thoughts in which

there is no reference failure. Such thoughts (thoughts in which there is reference failure) are just false; or, if the thought is desire-like, unsatisfied. (Searle 1983: 17) What preserves their content, despite the failure of reference, is that reference to particulars is effected descriptively.⁵² Two thoughts that differ from one another insofar as in one case, but not the other, there is reference failure can have the same contents because the content of the thoughts is not reference-dependent, but rather reference-determining.

My assessment of Searle has two parts. First, I examine phenomenological considerations that might count for or against Searle's account. Second, I present an argument in favor of my externalist account and show how it poses a challenge to Searle's internalism. It turns out that Searle's descriptivist strategy for dealing with cases of reference failure cannot be applied in certain contexts.

2.2.1: ASSESSING SEARLE: CONTENTS, COUNTERPARTS, AND PHENOMENOLOGY

Phenomenology is both a potential ally to the advocate of Searle's account and a potential enemy. Though Searle does not advert to this, it is an apparent virtue of his account that it affirms that counterparts possess identical contents. That is, I find it *prima facie* plausible that one cost of accepting the externalist view is that, according to it, counterparts differ in their contents. And so, it is *prima facie* a mark in favor of the comparative rationality of Searle's account that it holds that counterparts share contents. On the other hand, critics of Searle have argued that his account suffers from its postulation of contents that are not phenomenologically warranted. In this section, I will present and evaluate these two lines of thought.

⁵² The argument here is, of course, a reflection of Russell's from "On Denoting". (Russell 1905)

My evaluation of these lines of thought will depend on two connected ideas, intentionalism and the transparency of experience. An intentionalist characteristically claims that there is a tight connection between the content of a mental state and its phenomenal character. If visual experience is transparent, then introspecting visual experience reveals only features of the scene apparently revealed by visual experience. The claims I will be arguing for in this section are the following: The safest form of intentionalism is of little use to Searle. The strongest consideration in favor of intentionalism does not support the thesis that counterparts share contents. If experience were transparent then that might support a form of intentionalism, but not one that would make it reasonable to suppose that counterparts share contents. However, even if experience were transparent and this gave reason to suppose that counterparts must share contents, it is not clear that Searle's position is rendered all-things-considered reasonable by this. The reason is this: transparency can be wielded against Searle. However, since this line of attack against Searle is one that would also cause trouble for the externalist it does not make a difference to the comparative rationality of the positions.

One attractive feature of Searle's account is that it affirms, while the externalist theory denies, that counterparts of an experience share contents – and it does so while counting your experience in MBD1 as non-veridical. Remember that counterparts are experiences (inhabiting nearby possible worlds) that differ from one another with respect to their perceptual status, but do not differ with respect to their phenomenological character. The claim that counterparts share contents is sufficiently attractive that McGinn and Davies rely without argumentation on its truth in giving their theories of visual contents. But we must do better and seek the origins of the attractiveness of this

claim. If you join Davies in thinking that the notion of the content of experience is a phenomenological notion, then you'll have reason to accept that counterparts share contents – and thus reason to find Searle's view attractive.

Davies is committed to the principle that the counterparts of a perceptual experience share contents with it. This claim reflects his confidence that the theoretical *raison d'être* of a content of visual experience is to account for phenomenal differences and similarities among visual experiences. In the previous part of this chapter, I argued that cases like MBD1 cast doubt on the thought that visual content is a phenomenological matter. Indeed, they do – provided we set aside Searle's account. As argued previously, Searle's theory does not run into trouble with veridical misperception. It properly counts your experience in MBD1 as veridical. Consequently, Searle's account might preserve Davies' claim that visual content is a phenomenological matter by allowing that counterparts share contents.

What might someone mean in claiming that visual content is a phenomenological matter? One thing we might mean to signal by such a claim is the acceptance of a kind of reductionism about phenomenal character according to which the phenomenal character of an experience is constituted by its contents. Obviously, this thesis will have to be qualified somewhat. In general, it is unlikely that the phenomenology of a mental state can be entirely constituted by its content. This is because functionally different states might share contents, but differ in their phenomenology. So, for example, a conscious weak desire to eat a sandwich and a conscious strong desire to eat a sandwich plausibly share the same content, but differ in their phenomenology. For a more far-fetched example, suppose it were possible to form a judgment with the same content as a visual

experience. Probably the phenomenology of these states would differ: the judgment, for example, would have a word-y phenomenology, while the experience would not. Consider, moreover, the possibility of non-conscious experiences. While such experiences may well share contents with their conscious counterparts, they'll clearly differ in their phenomenal character. Part of the phenomenology of a mental state is determined not by its content, but by the kind of state it is – i.e. whether it is conscious, what sensory modality it embodies, what attitude it tokens, etc. Thus, the claim that the phenomenology of an experience is constituted by its content is presumably restricted. That is, it should be understood to apply only to those aspects of the phenomenology of experience that would not already be captured by a true theory of what makes a state a conscious visual experience.

Slightly more weakly, one might affirm that contents supervene on phenomenal character. This would mean, if true, that if two experiences are identical with respect to phenomenal character, then they are identical with respect to their contents. If we accept either the claim that phenomenal character is constituted by content or the claim that content supervenes on phenomenal character (call them the *Reductionist Theses*), then we're committed to thinking that counterpart experiences share contents.⁵³ Thus, to the extent these claims have warrant, we have reason to accept Searle's account.

The pressure here to accept that counterparts share contents derives from an acceptance of the thought that the contents of an experience are determined by its phenomenology together with the thought that counterparts are phenomenologically identical. Thus, one possibility for resisting the allure of Searle's account would be to

⁵³ A much weaker claim in this neighborhood is that phenomenology provides evidence about the contents of visual experiences.

deny that counterparts need be phenomenologically identical. Thus, one might accept not only externalism about the contents of visual experiences, but also externalism about their phenomenology; and, in so doing, diminish the appeal of Searle's account.⁵⁴

However, the prospects for this line of defense seem dim. In what follows, I'll use the experiences you have in MBD1 and MBD2 as representative examples. Call the experience you'd have in MBD1 *E1* and the experience you'd have in MBD2 *E2*. The thought that the phenomenal characters of the two counterpart experiences, *E1* and *E2*, are the same borrows some plausibility from two sources, each of which might be vulnerable. The first source of the thought that *E1* and *E2* share phenomenal characters is the thesis that the phenomenal character of mental states supervenes on the internal physical states of the subjects of those mental states. Except for what must be irrelevant differences, the states of your nervous system as you are subject to *E1* and *E2* are the same. So if we accept the supervenience thesis, then we ought to accept that *E1* and *E2* have the same phenomenal character. However, although this is a plausible view, I have yet to see a good argument for the supervenience thesis, I therefore propose to set this consideration aside.⁵⁵

The second source of plausibility for the thesis that the phenomenal characters of *E1* and *E2* are the same lies in the indistinguishability of *E1* and *E2*. It is probably true that *E1* and *E2* are indistinguishable by introspection. This (along with the claim that indistinguishable mental states must possess the same phenomenal character) yields the

⁵⁴ Mike Martin's disjunctivism, which I discussed in Chapter 1, is a form of externalism about phenomenal character.

⁵⁵ At any rate, the supervenience of phenomenology on internal physical states probably derives whatever warrant it possesses from considerations of indistinguishability, which I consider next. So an evaluation of this claim should give way to an evaluation of the role of indistinguishability in determining phenomenal character.

conclusion that E1 and E2 must have identical phenomenal characters. Thus, since there is reason to believe that E1 and E2 share phenomenal characters, there is reason to think that they must have identical veridicality conditions (if we accept the Reductionist Theses). Suppose, however, it were true that two states might be indistinguishable and yet differ with respect to their phenomenal character. If this were true, then we could defuse resistance to the thesis that E1 and E2 possess different veridicality conditions. We could accept that E1 and E2 are indistinguishable and we could accept the identification of phenomenal character and content, but could resist the conclusion that E1 and E2 must differ in their veridicality conditions.

However, this line of reasoning is problematic. The notion of phenomenal character is a philosophical invention and the claim that indistinguishable experiences might differ in their phenomenal character departs too far from the uses to which phenomenal character has been put. A bare denial that indistinguishable experiences share phenomenal characters cannot stand by itself. What I mean is this: if you would like to affirm the sentence “experiences might be indistinguishable but not share phenomenal characters”, you should replace the reference to phenomenal character with reference to whatever it is that you really care about (and what might normally be associated with phenomenal character). The position of someone who denies that indistinguishable experiences need be phenomenologically identical, is a bit like that of someone who denies that God is perfectly good. Since perfect goodness is criterial for being God, such a person should probably instead assert that the supremely-powerful world-creator is at least a little bit bad.⁵⁶

⁵⁶ Using indistinguishability as a criteria of type-identity for phenomenal characters is not philosophically unproblematic. In particular, doubts have been raised that distinguishability is transitive. If it is not

I'd rather not leave this as a bare terminological stipulation. Luckily, we can get some *very minimal* traction on the issue here. E1 and E2 are indistinguishable. Consequently, the way E1 seems and the way E2 seems are the same. If these states seem the same, then what it is like to be in them must be the same. If what it is like to be in E1 and what it is like to be in E2 are the same, then E1 and E2 are the same with respect to their phenomenal character.⁵⁷ Unless you are willing to abandon the idea that the phenomenal character of an experience is determined by what it is like to be subject to that experience, then you must accept that E1 and E2 possess the same phenomenal character.

It is highly probable that E1 and E2 share phenomenal characters and so now, if we want to resist those attractions of Searle's account that derive from its drawing a tight connection between phenomenology and content, we must look elsewhere. Of course, if the Reductionist Theses are false, then the advocate of Searle's account can find no support in them. Indeed, the Reductionist Theses are widely doubted.⁵⁸ The primary obstacle to accepting the Reductionist Theses comes from the conjunction of thoroughgoing externalism about the contents of mental states with internalism about phenomenal character. The content of an experience cannot supervene on its phenomenal character if the content does not supervene on internal states while the phenomenal character does. (Davies 1997: 323)

transitive, then it cannot be criterial for identity, since the identity relation is transitive. However, see Delia Graff's article, "Phenomenal Continua and the Sorites" for a review and resolution of these doubts. (Graff 2001)

⁵⁷ Eric Lormand's article, *The Explanatory Stopgap*, provides one way of making sense of the connection between the way a mental state seems and what it is like. (Lormand 2005)

⁵⁸ Andy Egan and James John's manuscript, "A Puzzle About Perception", provides a nice summary of the chief complaints against the identification of phenomenal character and representational content.

Given Davies' acceptance of the thesis that visual content is a phenomenological matter, he has some explaining to do. His solution to the problem is to suppose that the kind of supervenience postulated between content and phenomenal character and the kind of supervenience at issue in debates about content externalism are different.⁵⁹ In particular, they differ with respect to their modal strength. The supervenience at issue in debates about content externalism is what could be called *across-worlds supervenience* – which requires that necessarily “[if] x has ...property F in possible world w_1 , and y is a duplicate in w_2 of x (in w_1), then y has F in w_2 .” (Davies 1997: 313) The supervenience required for making good on the idea that visual content is a phenomenological matter could be called *within-a-world supervenience* for which it is necessary and sufficient that necessarily “[if] x has ...property F in possible world w , and y is a duplicate in w of x , then y has F in w .” (Davies 1997: 312) (Davies 1997: 323) Thus, on Davies account, even if content does not supervene across-worlds on internal states of creatures, the standard counterexamples (twin earths, inversions, and the like) do not establish that it doesn't supervene within-a-world.

Within-a-world supervenience of F on G , does not require that every possible thing that is identical with respect to being G be identical with respect to being F . It only requires that world-mates that are twins with respect to G be twins with respect to F . When Davies claims that within-a-world supervenience is sufficient for making it true that visual content is a phenomenal matter, he is surely overstating things a bit (which is lucky for the advocate of Searle's view because within-a-world supervenience is too attenuated to imply that counterparts share contents). At the very least, if visual content

⁵⁹ Davies borrows this strategy from McFetridge, who deploys it in defense of moral realism. (McFetridge 1985)

is a phenomenal matter, phenomenal twins in nearby possible worlds would be twins with respect to their contents. Whatever makes it true that F within-a-world supervenes on G should surely make it true that in nearby possible worlds G-twins would be F-twins. Consequently, we should expect that within-a-world supervenience is accompanied by a restricted version of across-worlds supervenience – according to which in some restricted domain of possible worlds (the nearby ones) F-twins are G-twins. Presumably, whatever substantive theory of the relation between content and phenomenal character (the theory that determines which contents supervene on which phenomenal characters) would determine which worlds count as nearby.^{60, 61}

The state of play, then, is this: the safest form of intentionalism is probably not committed to much more than the claim that counterparts share contents. Thus, it is not clear how it could serve as an independent motivation for the claim that counterparts share contents. In light of this, perhaps it would be best to proceed by seeking and evaluating the evidential roots of intentionalism. Why are the Reductionist Theses attractive – whether they are true or false – and can the roots of their attractiveness make it reasonable to accept Searle’s account of the contents of visual experiences?

Intentionalism may be motivated by engaging in an imaginative project in which one imagines changes to the introspected character of visual experiences and the

⁶⁰ Interestingly, Ludlow has argued (in effect) that the contents of propositional attitudes do not even within-a-world supervene on our internal states. (Ludlow 1995) Ludlow argues that changes in linguistic community – which can be a simple matter of switching from discourse with speakers of American English to discourse with speakers of British English – determine changes in the contents of our propositional attitudes. His arguments, however, cannot be naturally extended to cover the contents of visual experiences.

⁶¹ To be honest, I am not confident that it could ever make sense to postulate within-a-world supervenience without also being prepared to postulate at least some kind of across-worlds supervenience. What considerations would make it reasonable to suppose that world-mate twins that are the same with respect to G must be the same with respect to F, but not suppose that, at least for some range of possible worlds, cross-world twins that are the same with respect to G must be the same with respect to F? The only cases I can think of where it is reasonable to do this involve world-bound properties – e.g. being a denizen of w – and in these cases the obtaining of the within-a-world supervenience relation is trivial.

accompanying changes to their contents and vice-versa.⁶² Join me in fixing your attention on your current experience. I am looking at a round thing and my experience has a certain phenomenal character to it. Now I cannot imagine this very experience appearing to reveal a square thing rather than a round thing – while remaining the same with respect to its phenomenal character. If I try to imagine an experience that differs with respect to the shape of the thing I seem to see, I imagine an experience with a different feel to it. I imagine that were my experience representing different features, my subjectivity would be differently configured. And vice-versa: in imagining my visual subjectivity differently configured, I imagine things appearing to have different features. This exercise in imagination seems to make it reasonable to think that there is a tight connection – perhaps identity – between the phenomenal character of experience and its contents. Of course, by itself it does not succeed in establishing the truth any strong kind of reductionism. For example, we might have independent reason to think reductionism is false. For another, it is not entirely implausible that our powers of imagination are up to the task of capturing the kinds of cases that cast doubt on reductionisms.

However, suppose we set aside our doubts. Could the imaginative project above provide comfort to the friend of Searle's theory who wishes to find support for Searle's account in the fact that it predicts that counterparts share contents? I think not. We can accept considerations like the one above and yet affirm only that phenomenal character is constituted by the predicative content of visual experience. That is: if two visual

⁶² Another motivation is derived from the thought that the Reductionist Theses provide the "last best hope" for materialism about consciousness. However, as Speaks argues, there is little warrant for the Reductionist Theses to be found here. (Speaks 2007: 14) If phenomenal character and content are supposed to be identical, then content cannot be any more materialistically tractable than phenomenal character. One cannot, therefore, hope to get traction phenomenal character by reducing it to content. McGinn makes a similar argument in his article, "Consciousness and Content". (McGinn 1988) It is arguable that McGinn and Speaks commit the intensional fallacy here.

experiences are identical with respect to what properties they represent things as having, then they are identical with respect to their phenomenal character. Similarly, if two visual experiences are identical with respect to their phenomenal character, then they are identical with respect to their predicative content. However, even if this were true, it doesn't follow that phenomenally identical visual experiences are identical with respect to which particular objects they represent as bearing this or that feature. Hence, for all we've seen so far, we might accept a kind of intentionalism, while resisting the claim that the phenomenal type-identity of E1 and E2 commits us to their sharing the same veridicality conditions. They share the same predicative content (and this explains the appeal of supposing they share contents), but since they bear on different particulars, their contents differ. Put another way, we do not need Searle's account of the contents of experiences to account for the phenomenological similarities between E1 and E2 and so it cannot borrow any warrant from their phenomenological similarities.

The dialectic of the previous two paragraphs could be reproduced by considering another source of plausibility for the claim that visual phenomenology determines contents. Consider learning to taste a particular wine. How does one do this? I discover the content of my gustatory experience by attending to what it is like to taste the wine. That is, I can discover the content of my gustatory experience by engaging in the kind of introspective project which is supposed to bring the phenomenal character of experience to attention. I try to focus on various elements of the taste and thereby focus on various apparent elements of the wine – its particular fruitiness (is it a little bit foxy?), its acidity, its spiciness, its body, its tannins, its level of alcohol, its butteriness, etc. I can discover how my gustatory experience represents the wine by attending to what it is

like to taste the wine.⁶³ So, it might seem, there must be a determination relation between the phenomenology of the gustatory experience and its contents. However, this line of thought would establish at best only a limited kind of intentionalism – one according to which phenomenology is determined by some of the predicative contents of experiences.

One way in which phenomenology is both a friend and an enemy to Searle is that the transparency of experience is both a friend and an enemy to Searle. Let me first explain the use to which an advocate of Searle’s position might hope to put transparency. If visual experience is transparent, then introspecting visual experience reveals only features of the scene apparently revealed by visual experience. Transparency has been used to support the thesis that the phenomenology of experience is constituted by its content, which in turn would provide warrant for the claim that counterparts share contents.⁶⁴ This would provide just the support for Searle that I’ve been seeking in this section.

Earlier I argued that indistinguishable experiences share phenomenal characters. The rationale for this conclusion was that such experiences are the same with respect to what it is like to be subject to them. What it is like to be subject to an experience is determined by the face an experience reveals to introspection. Thus, it is plausible that the connection between indistinguishability and phenomenal character is grounded in the connection between phenomenal character and introspection.

⁶³ Of course, in so doing, I hope to discover two things: the objective features of the wine and the pleasures attainable in experiencing those objective features. Thus in learning to taste a wine, I am intrinsically interested in neither the contents of my gustatory experience nor its phenomenology. Attending to these things is instrumental.

⁶⁴ See Michael Tye’s article, “Representationalism and the Transparency of Experience”, for a recent application of the argument from transparency to the thesis that phenomenal character is constituted by content. (Tye 2002) Harman’s article, “The Intrinsic Quality of Experience”, is the contemporary source of this line of thought. (Harman 1990)

If what is revealed by introspection determines phenomenal character and introspection reveals content-involving facts, then content-involving facts determine phenomenal character. Suppose that visual experiences were transparent. If visual experience were transparent – that is, if introspection revealed only apparent features of the scene apparently revealed by visual experience – then a difference in phenomenology would imply a difference in content.

However, the line of thought in the previous paragraph is no utility to the advocate of Searle's position – who seeks a reason to think that counterparts share contents. One could accept that experience is transparent, and thereby receive warrant for thinking that differences in phenomenology entail differences in content, but not thereby receive warrant for thinking that a sameness in phenomenology entails a sameness in content. Consider (for an analogy) the claim that successfully listening to the oracle reveals only what Apollo is saying about Greece. This implies that if successfully listening to the oracle yields different experiences in two situations, then Apollo has said something different about Greece in those two situations. However, even if successfully listening to the oracle yields the same experiences in two situations, it is possible that Apollo has said something different about Greece in those two situations. Perhaps, though you hear *only* Apollo's speech about Greece by listening to the oracle, you do not hear *everything* he says about Greece in a situation by listening to the oracle. Maybe, for example, the oracle transmits only those of his utterances about Greece that are interesting to you. Similarly, transparency would provide some warrant for thinking that differences of phenomenology entail differences in contents, but, since it is possible that introspection doesn't reveal *all* the content-involving facts, it is possible to accept

transparency and deny that sameness of phenomenology entails sameness in contents. Some of the content of visual experience may well be phenomenologically silent.⁶⁵

While an advocate of Searle's account might affirm some variety of intentionalism (and thereby hope to find some warrant for Searle's account), he might also discover that accepting Searle's account is inconsistent with other aspects of the claim that the contents of visual experiences are a phenomenological matter. In particular, Searle's account might be found to conflict with substantive theses about the relation between phenomenology and content. One of the standard complaints about Searle's account is that the rich descriptive contents he postulates are not reflected in the phenomenology of visual experience. In particular, critics have found Searle's claim that perception-grounding relations are part of the contents of visual experiences to be implausible on phenomenological grounds.⁶⁶

The problem here is that transparency is also a threat to Searle's account. Suppose that the transparency thesis were true. When I seek to attend to what it is like to open my eyes in the presence of light it seems to me as if the phenomenal character of my visual experience is constituted by the scene in front of my eyes. This way of putting things is liable to misinterpretation. The phenomenology of visual experience is silent on its own constitution. It is not as if my eyes were telling me "your experience has such and such constitution" as they might tell me "that thing there is

⁶⁵ The argument of this paragraph and the previous one is adapted from Speaks article, "Transparency, Intentionalism, and the Nature of Perceptual Content". (Speaks 2007)

⁶⁶ See (Burge 1991) and (Armstrong 1991), for example. The phenomenological worry is not the only one in the neighborhood. For example, it can seem as if representing the dependence of visual appearances on objects of perception is beyond the representational capacities of the visual system. Such experiences must represent themselves and the dependence between themselves and states of affairs in the environment. However, the objection goes, this is probably beyond the powers of visual system. Soteriou puts the point in a rhetorical question: "What discriminatory abilities are left unexplained if one does not include the causal component in the content of visual experience?" (Soteriou 2007: 183)

purple”. Rather, if I were to explain to you what it is like to look around right now, I’d report something like this: it is like looking at a cluttered office, full of this and that. I’d report on the visual appearance of the scene in front of me. I’d report the felt character of my experience this way because in attempting to focus on my experience – in introspecting – I find myself attending to what appear to be distal objects and their features. Experience, as such, doesn’t show up as an apparent constituent of my visual phenomenology. When I open my eyes and look around, and simultaneously seek to attend to the phenomenology of my experience, I cannot single out something which appears to be an experience.⁶⁷ By hypothesis, the same is true of all of us. The problem for Searle is supposed to be this: since experience doesn’t show up as an apparent constituent of visual phenomenology, its causal dependence cannot show up as an apparent constituent of visual phenomenology.

If you think that only what is an apparent constituent of visual phenomenology can be a constituent of content and you accept the transparency of experience, then you’ll have reason to reject Searle’s theory. As is not surprising, Searle himself denies that perception-grounding relations need be reflected in the phenomenology of visual experience.⁶⁸ (Searle 2004: 326) (Searle 1991: 236) On Searle’s view, some of the content of a visual experience contributes to the phenomenology of being subject to it. Some of the content does not. That aspect of the content of experience that determines

⁶⁷ It is important to note that this is compatible with the claim that in fact when one introspects one’s experience and its features are in fact objects of introspective awareness. However, if it is true that we are aware of our experience and its features, then we are not aware of them *as* our experience and its features. That is, it is possible that the transparency of experience is an *illusion*.

⁶⁸ There is, he claims, phenomenological support for the thesis that some *non-visual* experiences actually have such contents. He reports that in cases of tactile perception, in cases of sensing an object’s contact with his body, for example, it seems to him as if there is a sensation and at the same time that the sensation is caused by something. “I feel the object as causing me to have this sensation.” Thus it is possible, he thinks, for sensuous awareness to reflect perception-grounding contents in its phenomenology. (Searle 1991: 236)

how things look contributes to the phenomenology of experience. That aspect of the content of experience that represents the perception-grounding relations does not normally contribute to the phenomenology of experience. Thus he is committed to denying the principle that the contents of visual experiences are limited to what is reflected in phenomenology.

Some reason for suspecting that visual experiences represent what is reflected in their phenomenal character might be found in my earlier discussion of learning to taste a wine. One discovers the appearance of a thing (how it is revealed in one's experience) by introspection, by attending to its phenomenal character. This procedure would not make sense unless the content of one's experience was manifest in its phenomenal character. However, as should be clear from earlier argumentation, this doesn't entail that all of the content of visual experience is manifest in its phenomenal character. That is, it doesn't support the hypothesis that visual experiences represent *only* what is reflected in their phenomenal character.

However, it is not clear that we need to decide whether visual experiences must represent only what is reflected in their phenomenal character. The reason is this: if Searle suffers from postulating elements of the content of a visual experience that are not reflected in visual phenomenology, then so does the externalist. According to the externalism I favor, there can be experiences that share phenomenal character but differ with respect to their contents. Some differences in content make no phenomenal difference on the externalist view. Like Searle, the externalist is committed to affirming that there is more to the content of experience than is reflected in the phenomenology of

experience (on the assumption that being reflected in the phenomenology experience requires being a phenomenological difference-maker).

It is time, then, to take stock. We have seen that there is some reason to think that the predicative contents of visual experiences supervene on phenomenology. However, this would not establish that counterparts must share contents since counterparts might share predicative contents, while differing with respect to their bearing on particulars. We have seen that transparency might warrant only a weak version, according to which phenomenological differences entail differences in content. What the friend of Searle would want here is a stronger thesis according to which phenomenological type-identity entails content-type identity. Additionally, while a consideration of transparency, together with the thesis that visual experiences must represent only what is reflected in their phenomenal character, would cause trouble for Searle, it would cause similar trouble for the externalist.

It seems that phenomenology is neutral with respect to the comparative rationality of Searle's account and my externalist account. Thus the advocate of Searle's account will find neither consolation nor discomfort by its consideration. While it seemed at first glance that it was an advantage for Searle's account over the externalist account that his account allows counterparts to share contents, it is not. Contrary to Davies and McGinn, it does not seem that the notion of a visual content is a phenomenological notion. Phenomenology plays some role in a theory of visual content (by giving evidence about the predicative contents of visual experiences, for example), but there are evidently other constraints on the postulation of contents for visual experiences, other reasons to postulate visual contents.

For example, as we have already seen in my assessment of the veridicality of MBD1, one function of visual contents is to account for our judgments of veridicality. Thus, we have intuitive reason to suppose that your experience in MBD1 was non-veridical. This constrains the contents of your experience. It tells us that whatever the content of your experience is, in this case, it had better depend for its veridicality on the states of the seen objects. Your experience must represent them. Thus, either it represents them directly – as I urge – or it must represent them indirectly (by representing the obtaining of conditions which they in fact satisfy) – as Searle urges.

2.2.2: ASSESSING SEARLE: VISUAL CONTENTS AND THE MENTAL ENVIRONMENT

Some constraints on the contents of visual experiences can be argued to come from the relations between visual experiences and other mental states. Accordingly, we can expect to find reasons to reject or accept Searle's account by looking into the relations between visual experiences and other mental states. In this section, I'll look at two such cases. First, I'll critique an argument Searle gives for his account that rests on the introspectibly available differences between visual experience and visual imagination. Second, I'll present an argument that casts my externalism about visual contents in a better light than Searle's account. The argument will focus on a difference the obtaining of the seeing relation makes to the respects in which our actions are intentional.

2.2.2.1: THE MENTAL ENVIRONMENT: DIFFERENCES AMONG ATTITUDES/SENSORY MODALITIES

Searle thinks that some warrant may be found for his theory in the phenomenological differences between exercises of the visual imagination and in exercises of our capacity to see:

Suppose we had the capacity to form visual images as vivid as our present visual experiences. Now imagine the difference between forming such an image of the front of one's house as a voluntary action, and actually seeing one's house. In each case, the purely visual content is equally vivid, so what would account for the difference? The voluntarily formed images we would experience as caused by us, the visual experience of the house we would experience as caused by something independent of us. (Searle 1983: 124)

As he puts it, substantially more strongly, in a later writing:

...you can indirectly bring the causal conditions of satisfaction to consciousness. In the case of perception there is an experienced contrast between the voluntary character of visual imagination, where my intention causes the visual image, and my actual visual perception, where I experience the visual experience as caused by objects in the external world. (Searle 2004: 326)

The argument here is a kind of inference to the best explanation.⁶⁹ Given that being in the kind of mental state that issues from acts of visual imagination differs phenomenologically from being in the kind of mental state that results from exercises of the capacity to see, how best to explain the difference? Searle thinks that if exercises of the capacity for seeing issued in experiences that represent perception-grounding relations and the same is not true of exercises of the capacity for visual imagination, then we could make sense of the phenomenological difference between seeing and visually imagining. The phenomenology of these mental states differs because perception-

⁶⁹ It is a bit troubling that in the first quotation Searle speaks of the "purely visual content" of a visual experience. Does this suggest that Searle thinks that the causal condition is not part of the purely visual content of a visual experience? What distinction does Searle have in mind in suggesting a difference between purely visual content and content that is not purely visual? How would such a distinction affect the dialectical situation?

grounding relations are part of the contents of visual experiences, but not part of the contents of the states that issue from the exercise of the visual imagination.

However, as I will now argue, there is reason to doubt that the attitude or sensory mode of a mental state is contained in its contents. If so, we should expect that our knowledge of the difference between exercises of the capacity for seeing and exercises of the capacity for visual imagination would not depend on differences in the contents of these mental states. The problem comes out most clearly in the case of desires. If, as seems plausible, the content of a desire is exactly what is desired (and therefore what structures the actions and plans motivated by the desire), then the fact that one is desirous cannot be supposed to be generally part of the content of a desire. Consider, for an example, the desire that I have a sandwich for lunch. If the fact that one is desirous were part of the content of the desire, then the desire would be perspicuously represented by something like the following:

Desire (I am desiring that I eat a pulled pork sandwich).

But this would entail that those desires we took to be first-order desires are really second-order desires – desires to desire. If the attitude – desiring – were itself part of the content of a desire, then my desire for a sandwich would, in addition to inclining me to taking the steps required to eat a sandwich, would incline me to maintain my desire for the sandwich. But it is clearly false that no desire is a first-order desire. We can find ourselves alienated from our own desires. For example, I can find myself desiring to eat a pulled pork sandwich, but, at the same time, wish I could be rid of my desire (e.g. out of

sympathy for the plight of pigs). I might – consistent with my desiring to eat a pulled pork sandwich – not desire that I desire to eat a pulled pork sandwich.⁷⁰

It might be objected that desires are special, and that mental states with different directions of fit work differently. However, consider hypothesizing, which has a belief-like different direction of fit rather than a desire-like direction of fit. You can ordinarily tell whether you are entertaining a proposition in a hypothesizing way rather than in a believing way. It would be very strange if the fact that one is hypothesizing were part of the content of one's hypothesis was what explained this. Consider for example hypothesizing that visual experience is transparent. If the fact one is hypothesizing were part of the content of one's hypothesizing state, then it could be perspicuously represented as follows:

Hypothesize (I am hypothesizing that visual experience is transparent)

If this were an explicit (and true) representation of the content of a hypothesis that visual experience is transparent, then we'd have trouble. For example, we use this kind of hypothesis to see what would be true if visual experience were transparent. However, if the content of a hypothesis were that one hypothesizes such-and-such, then we could not use such a mental state to see what would happen if visual experience were transparent. For example, since one might hypothesize what is false, the hypothesis that one hypothesizes that *p* does not license the hypothetical conclusion that *p*.

We should expect that a person's knowledge of what kind of mental state he is in can be explained by something other than the content of the mental state. Thus, when Searle attempts to explain the introspectibly available difference between the exercise of

⁷⁰ I expect there might be regress problems in the offing here as well.

the capacity for visual imagination and the exercise of the capacity for seeing by postulating differences in the contents of the mental states involved we should doubt him. We often enough know when we are desiring or hypothesizing – though our desiring is not part of the content of our desire and our hypothesizing is not part of the content of our hypothesizing. And I don't see any compelling reason to suppose that things work differently in the case of knowing we are seeing rather than visualizing.

2.2.2.2: THE MENTAL ENVIRONMENT: PERCEPTION, THE PROPOSITIONAL ATTITUDES, AND MENTAL ANAPHORA

In this section, I'll present an argument that casts my externalism about visual contents in a better light than Searle's account. The argument will focus on a difference the obtaining of the seeing relation makes to the respects in which our actions are intentional. There is some warrant for the conclusion that whether the seeing relation obtains makes a difference to the respects in which our actions are intentional. The most elegant account of this difference-making holds that this difference-making begins with visual contents and spreads from there to the contents of our beliefs, desires, and decisions. The problem for Searle is that his strategy for preventing the spread of proto-semantic difference-making causes trouble for his internalism.

I'd like to begin by thinking about a twin story taken from Sean Crawford's article, "In Defense of Object-Dependent Thoughts" – call this tale *C1*. (Crawford 1998) Like many such narratives, *C1* gives us an opportunity to reflect on the similarities and differences between twins. Consider, then, Ralph and Twin Ralph (*TR* henceforth):

TR and Ralph live in the same world, share histories that are exactly intrinsically similar, share roughly the same propositional attitude psychology, each is a molecular duplicate of the other, etc. Ralph and TR are located in exactly similar

rooms and in front of each there is a table and a glass of lemonade. Ralph sees the glass of lemonade on the table. But, TR is the victim of a veridical hallucination. That is, TR is subject to experience just like Ralph's (to the extent that this is compatible with not actually seeing what Ralph sees) and his environment is exactly similar to Ralph's environment.

Unlike Ralph, TR fails to see a glass of lemonade. For some reason, TR's experience has the kind of causal history that is insufficient to make it true that he sees the glass in front of him. However, if there were any cases of veridical hallucination, TR's would be suffering from one of them. Just as with Ralph, things are as they look to TR.

Both of them reach for their glasses of lemonade. Or, more carefully, both of them stretch out their hands, grasping thus-and-so for egocentrically identical regions of space. Their action is the upshot of a train of thought they might both express in the following terms: "I am thirsty for lemonade. That is a glass of lemonade. Grasping that is what I must do in order to satisfy my thirst. OK: I shall grasp that and drink it."

Obviously, TR and Ralph are rather similar psychologically. They are, in some sense, behaving the same way. They would use the same words to explain their behavior. If TR and Ralph have ulcers, we could accuse each of being foolish ("As you well know: lemonade will make your ulcer worse!") Neither of them is in any position to distinguish his own state from his twin's. Etc.

However, there is reason to suppose that Ralph and TR are interestingly different psychologically. Crawford claims, on the authority of intuition, that Ralph and TR differ with respect to their agency. (Crawford 1998: 204) Ralph's action is intentional_a with respect to which glass he reaches for.⁷¹ As I'll explain in a moment, that it is intentional_a in this respect places constraints on the contents of Ralph's mental states. TR's motion is not similarly intentional_a: there is no glass on which TR intends_a to act – even though he

⁷¹ In this section, the reader might be confused by an ambiguity in the word 'intentional'. Until this point, I've been using 'intentional' as a term that refers to the aspects of content. However, here we are interested in differences between Ralph and TR's agency – difference with respect to their intentions. Thus, for the duration of this section, I'll subscript 'intention' (and words built on its stem) with an 'a' when I mean to talk about aspects of agency and I'll subscript it (and words built on its stem) with a 'c', when I mean to talk about aspects of content.

grasps for the glass in front of him. Thus, unlike TR, Ralph can and does act intentionally_a on the glass in front of him – that it is the glass in front of him on which he acts is one respect in which his action is intentional_a. As Crawford puts it, “he [that is, TR] did not intend to pick up *that* glass.” (Crawford 1998: 206) The difference between them is a difference with respect to whether their actions are intentional_a in relation to some particular object.

Before I move on to a discussion of the way in which these postulated agential differences justify claiming that Ralph and TR differ in the contents of their thoughts, it is worth pausing to add some support to Crawford’s claim that Ralph and TR differ with respect to their agency. Is it really true that Ralph does, while TR does not, act intentionally_a on the glass in front of him? My judgment shifts back and forth on this question. Crawford calls TR, *lucky Ralph*. Perhaps, given that TR is lucky, this justifies his claim that Ralph and TR differ with respect to their agency.

How is TR lucky and could TR’s luckiness be an obstacle to the particularity of his action?⁷² My account of the perception relation would have the implication that TR is lucky, given that his experience is hallucinatory, should he succeed in grasping the cup. His experience is not controlled by the cup and its features. Even if there were a reverse censor waiting offstage, the process by which TR came to have the experience he has (and consequently the propositional attitudes he forms) is not one that reliably produces

⁷² It is commonly supposed in discussions of moral responsibility and indeterminism that luck is a problem for one’s being morally responsible. See Galen Strawson’s article, “The Impossibility of Moral Responsibility”, for an example of this. (Strawson 1994) For example, if you reach a decision by luck (as a result of random fluctuations in one’s neurons, say), then this plausibly undermines your moral responsibility for the resulting action. However, we cannot model the case of TR on this kind of case. For in a case in which your action is controlled by a decision (which was reached by luck), your action might nonetheless be intentional in every respect in which it would be were it the consequence of a decision that was not reached by luck. This kind of luck would undermine moral responsibility without necessarily changing the intentional_a facts.

accurate experiences. There could have been, for some examples, a crucible of molten lead, a snarling pit bull, or nothing at all in front of him and he would still have made decisions like those he actually forms. Thus, that TR decides to reach for the cup of lemonade when there is one in front of him is a lucky circumstance, given that he was hallucinating. That TR's decision is likely to result in his sipping lemonade is similarly lucky. It is amazingly lucky that his hand is configured just so, and his arm will extend just enough (but not too far!) to grasp the glass of lemonade, given that he cannot see it. TR is lucky because his actions will lead to satisfaction, even given that he cannot see what would normally be required for success.

Why would TR's being lucky in this way undermine the claim that his action is intentional_a with respect to its being *that* glass on which he acts? Suppose that TR was sitting in front of many glasses and because the room in which he sits is in total darkness he can see none of them. If he should reach out and grasp a glass (without knocking it or its neighbors over) his success would be lucky. Moreover, it is clear that his reaching out and grasping that very glass would be lucky. Accordingly, I think we would not say that, in this case, his action was intentional_a with respect to its being that glass he picks.⁷³ We must remind ourselves that the case of C1 is like this case. In C1, TR might as well be sitting in a darkened room given that he cannot see the glass in front of him. Accordingly, he is lucky in whatever success he meets and, therefore, does not act intentionally_a on that glass (the one in front of him).

Why do the agential differences – supposing there are such – between Ralph and TR justify claiming that they differ psychologically? Differences in their capacities for

⁷³ Suppose that his reaching out and grasping was not preceded by some clever bout of intellection by which he derived from first principles the layout of the darkened room.

the exercise of agency suggest that Ralph and TR are thinking thoughts with different kinds of content. If Ralph intends to act on a particular glass of lemonade (rather than another similar glass or nothing at all), then there is probably something about the content of his intention_a that is sufficient to make it true that it is *that* particular glass of lemonade (rather than another similar glass or nothing at all) on which he acts. It appears that due to differences in the perceptual facts, Ralph possesses intentional_c capacities TR lacks. Ralph's possession (and TR's lack) of that kind of intentionality_c is what explains their different capacities for intentional_a action.

This connection between the content of intention_a and the respects in which an action is intentional_a is the one of the basic insights behind causal theories of intentional_a action. I might intend_a to move my arm, and in doing so destroy Tokyo. As it happens, the movement of my arm awakens Godzilla, who proceeds to level the city of Tokyo. If someone were inclined to blame me for this unhappy event, I might protest that while it is true that I awoke the fearsome beast and thereby destroyed Tokyo, it is false that I did this intentionally_a. The assignment of blame is made problematic (though not necessarily defeated) by the fact that the content of my intention_a was simply that I move my arm. In virtue of being caused by an intention_a with this content, my action was intentional_a in respect to being an arm movement, but not in respect to being a destruction of Tokyo.

If thoughts with the relevant contents are sufficient to make it true that Ralph's grasping of the glass is relevantly intentional_a, then it had better be the case that TR is not thinking thoughts with similar content. If TR were, then his action would also be intentional_a in the relevant respects. But, we're supposing, his action is not intentional_a in

the relevant respects, so it is likely that he is not thinking thoughts with contents like Ralph's.

It appears, then, that the perceptual facts make a difference to the contents of one's intentions_a. That Ralph is seeing and TR is not apparently makes a difference to the contents of those thoughts that determine the respects in which their actions are intentional_a. Thus, it would appear, the perceptual facts make a proto-semantic difference with respect to those thoughts that determine the respects in which their actions are intentional_a. The most elegant account of this difference-making holds that this difference-making begins with visual contents and spreads from there to the contents of our beliefs, desires, and decisions. Call this view of the matter the *Maximal View*. In order that the case of Ralph and TR support my externalism about the contents of visual experience, it had better support the Maximal View. However, it would be possible for a theorist to accept more modest accounts – call them the *Minimal View* and the *Intermediate View*.

The Minimal View holds that the perceptual facts make a difference only to the contents of their decisions and intentions. The Minimal View has the virtue of evidential parsimony. It accepts no more than might be required in order to account for C1. The Intermediate View holds that the perceptual facts make a difference not only to the contents of decisions and intentions, but also to the contents of the propositional attitudes generally. However, the perceptual facts make a difference only to the contents of the propositional attitudes and not to the contents of visual experiences.

Though the Intermediate View is less evidentially parsimonious, it possesses compensating charms. We should expect that the contents of our decisions and intentions

match the contents of our other propositional attitudes. The reason is this: our intentions_a and decisions are characteristically related inferentially to our beliefs and desires. Decisions and intentions live in a deliberative environment in which they are arrived at via inferential connections to our beliefs and desires. So for example, should I feel a desire to eat a pickle and judge that one is available in my refrigerator, I might decide (on that basis) that I shall get a pickle from my refrigerator and eat it. One theoretical role for the contents of mental states is to make sense of these transitions by rationalizing them. These transitions are intelligible because there is an overlap between the contents of the transitioned-between states. For this reason, it would be surprising if Ralph and TR's decisions bore on particulars in one way, while their desires and beliefs (the ones to which their decisions are related in deliberation) bore on particulars in a different way. Thus we should expect that if the perceptual facts make a difference to the contents of their decisions and intentions, they also make a difference to the contents of their propositional attitudes generally.

The Maximal View holds that the perceptual facts make a difference all the way down. They make a difference not only to the contents of their propositional attitudes (as required by the Intermediate View), but also to the contents of their visual experiences. It is even less evidentially parsimonious than the Intermediate View, but is more elegant. We already have some reason to think that the contents of visual experiences depend on the perceptual facts and so those reasons – together with the reasoning that takes us from the Minimal View to the Intermediate View – would recommend the Maximal View. The externalist result of considering MBD1 coheres nicely with the externalist result of considering C1 – thus each borrows warrant from the other.

Additionally, if we affirm the Intermediate View and deny the Maximal View, then we have an awkward gap. What I mean is this: on the Intermediate View, the contents of the propositional attitudes depend on the perceptual facts, but the contents of visual experiences do not. This raises the question: why should it be this way? We have already considered and rejected phenomenological considerations that might make it reasonable to deny that the contents of visual experiences depend on the perceptual facts.

You might think that the non-conceptual character of visual experience would make this gap less awkward by making it predictable. It is plausible that sensory experience represents the environment more fine-grainedly than human judgment. This has been used to argue that sensory experiences possess what are called *non-conceptual contents*.⁷⁴ Sensory experience is fine-grained, which means that “visual experiences represent the world with a determinacy of detail that goes beyond the concepts possessed by the subjects of those experiences.” (Tye 2006: 519) The usual example is the representation of color in visual experience. My experience-producing system has the capacity to represent a great many shades. Some information as to how many shades of color my experience can represent can be found in the number of shades computer designers design computers to display. Photographic quality pictorial representations require the use of a format that allows for the representation of millions of colors (as with 24-bit displays, for example). However, it is plausible that I do not possess millions of distinct color concepts. In forming non-demonstrative judgments about the colors of things, I can deploy only a small handful of color concepts. Thus, it is plausible that visual experiences represent the world in a determinacy of detail that outruns the capacity

⁷⁴ See Tye’s article, “Nonconceptual Content, Richness, and Fineness of Grain”, for a discussion of the issues. (Tye 2006)

of my judgment-forming system. Evidently, the representational resources of the visual system differ from the representational resources of the cognitive system.

However, there is no reason here to suppose that these differences in the representation resources of each system would mean that the contents of the propositional attitudes depend on the perceptual facts while the contents of visual experiences do not. The reasons for supposing that these experiences possess non-conceptual content are reasons concerning only the predicative contents of these states. There is reason to believe that what explains the fittingness of states of the cognitive system for representing properties differs from what explains the fittingness of states of the visual system for representing properties. Thus we have some reason for thinking that the predicative contents of states of the two systems might differ in their proto-semantics. However this does not provide reason for thinking that the bearing of states of one system on particulars might depend on the perceptual facts, while the bearing of the states of the other system does not.

The Maximal View affirms that the perceptual facts make a difference to the contents of visual states and the contents. They do so proto-semantically – by fixing those contents. What view might Searle take of the matter? Searle is committed to thinking that Ralph and TR share intentions with the same contents. Searle's reason for thinking this stems from his generalized internalism – which in turn depends on his view of reference failure.

Accepting Searle's commitment to the idea that the contents of intentions do not differ just because the perceptual facts are different, explaining the fact that Ralph's action is intentional_a in ways that TR's is not must not proceed by invoking a difference

in the contents of their intentions_a.⁷⁵ How would such a explanation proceed? Let's imagine Searle agreeing that Ralph's intention_a is a demonstrative intention_a. If Ralph's intention_a is demonstrative, then, on Searle's account, so is TR's.

Here is Searle's representation of the content expressed by an utterance of the sentence 'That man is drunk':

((there is a man, x, there, and the fact that x is there is causing this visual experience) and x is the man visually experienced at the time of this utterance and x is drunk) (Searle 1983: 227)

Can we transpose this account of the content of a demonstrative statement to an account of the content of a demonstrative intention_a? Let's see.

Consider, then, the intention_a that *that* man be drunk. Suppose, for some scene setting, that I'm at a dinner and trying to leave a potential employer with a good feeling about me. Perhaps it would be a good idea to keep the drinks flowing. What, then, might a Searle-style analysis of the content of such an intention_a look like? Begin with a preliminary statement of my intention_a: I will make it true that *that* man gets drunk. Searle's account of how to deal with perceptual demonstratives is given by his representation of the content of 'That man is drunk.' Searle thinks that this statement has, as part of its satisfaction conditions, the obtaining of reflexive causal relations. Let's, then, apply this same move to the case of the intention to make it true that *that* man gets drunk.

I will make it true that: ((there is a man, x, there, and the fact that x is there is causing this visual experience) and x is the man visually experienced at the time of this intention and x is drunk)⁷⁶

⁷⁵ Remember that we have set aside differences that involve their being different people and being subject to different token mental states.

⁷⁶ There is some difficulty here about tense. I have imagined forming the intention_a to get that man drunk at t1, while I'm looking at him. The problem I am about to pose for Searle is one having to do with thinking that the identifying information is part of the content of the intention_a. Evidence against thinking

Here we can start to see the problems with the Searle-style move. The main difficulty is this: not every condition required by the proposed content of the intention_a is such that I intend_a to make it true. For example, I do not, in intending_a that *that* man gets drunk, thereby become inclined to take any steps to control the location of the potential employer. All of my efforts, presumably, will focus on keeping his glass full of delicious, inebriating liquids. It seems that my actions are not at all focused on making true propositions about my perception-grounding relations to the man I wish to inebriate.

We should defend Searle's view by claiming that it is a mistake to include the reflexive causal conditions – i.e. those involving visual experiences and such – in the content of the intention_a. Perhaps we should think the following: insofar as intentions_a bear on particulars they are accompanied by beliefs that represent the obtaining of Searle's proposed reflexive causal conditions. When I intend_a that *that* man get drunk, I believe, let's say, that there is a man there and he is the cause of my current experience and intend_a of him that he become drunk. My intention_a borrows its bearing on that man from the beliefs that accompany it. The bearing of intentions_a on particulars is, so to speak, anaphoric on our belief-like states.⁷⁷

this comes from the fact that none of the action or planning consequent to forming my intention_a would involve making it true that I bear perception-grounding relations to the man. One difficulty for this argument is that later – as I am acting or planning – the identifying descriptive content will be most perspicuously expressed in past tense. A defender of Searle might note that the explanation for why none of my action or planning focuses on making these things true is that they are already true and so beyond my capacity to make true. This would explain my inaction with respect to the obtaining of identifying facts – and does so without denying that these facts are part of the content of the intention. However, such a proposal runs afoul of the thesis that the content of intentions_a should be limited to what is to be made true. Intentions_a have, on Searle's theory, a characteristic direction of fit. If such-and-such state of affairs is not represented as being to be made true, then it is part of the content of an intention_a, but rather part of the content of a different kind of state.

⁷⁷ There is reason to believe that Searle himself would favor this analysis. (Searle 1983: 216)

It is worth mentioning that the dialectic of the previous two paragraphs will generalize for a wide variety of desire-like states involving existential commitment. Consider, for example, the wish that the people at my party leave. This desire expresses an existential commitment to its being true that there are some people at my party. However, it would be a mistake to include this existential commitment in the content of my desire. After all, it would not necessarily satisfy my wish if there were people at my party at the time of my wish and they subsequently left. It is compatible with wishing that the people at my party leave that I also wish they had never arrived in the first place. Thus, on any account, it will often be the case that desires we ascribe by using existentially committed language in the clause specifying the content of the desire will have to be understood to be anaphoric on accompanying belief-like states.⁷⁸ Descriptive desire-like states will have to be understood as anaphoric on accompanying belief-like states. In Searle's case, since he gives demonstratives a descriptive analysis, it will turn out that every demonstrative desire-like state is such that its bearing on a particular is anaphoric on accompanying beliefs.⁷⁹

In the case of demonstrative intentions, the beliefs in question will presumably be demonstrative. Since demonstrative beliefs get a descriptive analysis on Searle's account, we can perspicuously represent Ralph and TR's mental state (insofar as they intend demonstratively) as follows:

Believe (something, x, is S and the fact that x is S is causing this visual experience and x is a glass of lemonade.)⁸⁰

⁷⁸ It is beyond the scope of the dissertation to give an account of the conditions under which we'll have to postulate anaphoric dependence for desires.

⁷⁹ If there were reason to deny that demonstrative desires (and desire-like states) were anaphoric on accompanying beliefs, then we'd have reason to reject Searle's account. Such a case might proceed by showing that there are desire-like states that are not accompanied by the right kind of beliefs. I cannot think of any such cases, however.

⁸⁰ 'S' expresses the predicative content of the experience.

Intend (that I grasp it)⁸¹

The bearing of the intention on the glass of lemonade picks up its reference from the belief represented above it. Since the bearing of Ralph and TR's intentions on particulars is anaphoric on the Searle-style beliefs that accompany them, the fact that TR's belief is false has implications for the content of his intention. It looks like we shall have to suppose that TR's intention suffers from reference failure. Since there is nothing that satisfies the condition (that there be something that is S and...) required by the content of TR's belief, the content of his belief cannot supply reference for his intention_a.

That TR's intention suffers from reference failure is not a surprise. I take it that this is common ground between the externalist and Searle. Indeed, it is reference failure that probably accounts for the difference in agency between Ralph and TR. TR's action is not intentional_a with respect to which glass he grasps for because his intention fails to refer to the glass. Consequently, it is not part of the satisfaction conditions of his intention that he grasp *that* glass. The problem for Searle is this: If TR's desire suffers from reference failure, then even a Searle-style account cannot hold on to its affirmation that Ralph and TR's intentions share the same contents. Remember, Searle's strategy for preserving sameness of content (despite reference failure) depends on postulating particularity-determining descriptive contents. However, we have seen that desires do not really have particularity-determining descriptive contents. That is, when we ascribe a desire-like state using a particularity-determining description in the content-specifying clause of the desire ascription, we are indicating a state of affairs in which a desire is

⁸¹ I use 'it' in preference to 'x', since the use of 'x' would suggest that the 'that I grasp x' is an open sentence or that 'x' is a variable bound by the quantifier in the representation of the belief. I want to make no such suggestion.

accompanied by a descriptive belief and takes its reference from that belief. Thus, Searle cannot apply the descriptivist strategy for dealing with reference-failure to the case of desire-like states.⁸²

But now, if we reject the claim that Ralph and TR share intentions_a, we have no reason to prefer Searle's view over its externalist rival. Whatever account he gives of the content of TR's intention it needn't be a descriptivist account (because it cannot be one). It must be possible to address the problem of reference failure without postulating descriptive contents. I see no reason to expect a move that will work in the case of desire-like states cannot be applied across the board to other content-bearing mental states.

CONCLUDING REMARKS

We have seen that there is no reason to be found in the connection between phenomenology and visual content to prefer Searle's account over my externalist analysis of MBD1. However, it seems that there is no reason to reject Searle's account in light of the connections between phenomenology and visual content. While Searle has the resources to account for the case of Ralph and his twin, he must hold that the Ralphs' desire-like states are anaphoric on accompanying beliefs. This has the consequence that TR's intentions suffer from a failure of reference that cannot be accounted for using descriptivist resources. Thus, one of the chief reasons for being attracted to Searle's internalism is undermined by the anaphoric relation he must posit to account for the

⁸² What account should be given of the content of TR's intention? I will not give one here. Since my present goal is to seek differences between my view and Searle's that cast Searle's view in worse light than my own, I need not give an account of the content of TR's intention at this point. Searle and I both stand in need of an account of the content of TR's intention.

Ralphs. Additionally, there is an elegant externalist account that would also account for the Ralphs. At any rate, we have seen no reason to accept Searle's account that would outweigh the intuitive plausibility of the claim that veridical hallucination is possible.

BIBLIOGRAPHY

- Alston, W. (1995) 'How to Think about Reliability'. *Philosophical Topics* 23: 1-29.
- Alston, W. (1999) 'Back to the Theory of Appearing'. *Philosophical Perspectives* 13: 181-203.
- Armstrong, D. (1991) 'Intentionality, Perception, and Causality: Reflections on John Searle's Intentionality'. In *John Searle and His Critics*, eds. E. LePore and R. Van Gulick. Oxford: Basil Blackwell.
- Burge, N. (1991) 'Vision and Intentional Content' In *John Searle and His Critics*, eds. E. LePore and R. Van Gulick. Oxford: Basil Blackwell.
- Child, T. W. (1994) *Causality, Interpretation and the Mind*. Oxford: Clarendon Press.
- Comesaña, J. (2005) 'Justified vs. Warranted Perceptual Belief: Resisting Disjunctivism'. *Philosophy and Phenomenological Research* 71: 367-383.
- Comesaña, J. (2006) 'A Well-founded Solution to the Generality Problem'. *Philosophical Studies* 129: 27-47.
- Conee, E. and Feldman, R. (1998) 'The Generality Problem for Reliabilism'. *Philosophical Studies* 89: 1-29.
- Crawford, S. (1998) 'In Defence of Object-Dependent Thoughts'. *Proceedings of the Aristotelian Society* 98: 201-210.
- Dancy, J. (1995) 'Arguments from Illusion'. *The Philosophical Quarterly* 45: 421-438.
- Daredevil*. (2003) Dir. Mark Steven Johnson. Twentieth Century-Fox Film Corporation.
- David Widerker and Michael McKenna, eds. (2002) *Moral Responsibility and Alternative Possibilities*. Aldershot, UK: Ashgate Publishing.
- Davidson, D. (1980) *Essays on Actions and Events*. Oxford: Clarendon Press.
- Davies, M. (1983) 'Function in Perception'. *Australasian Journal of Philosophy* 61: 409-26.
- Davies, M. (1992) 'Perceptual Content and Local Supervenience'. *Proceedings of the Aristotelian Society* 92: 21-45.
- Davies, M. (1997) 'Externalism and Experience'. In *The Nature of Consciousness*, eds. N. Block, O. Flanagan, and G. Guzeldere. Cambridge, Mass: MIT Press.

- Dummett, M. A. E. (1979) 'Common Sense and Physics'. In *Perception and Identity: Essays Presented to A. J. Ayer with his Replies*, ed. G. F. MacDonald. Ithaca, NY: Cornell University Press.
- Egan, A. and John, J. (MS.) 'A Puzzle About Perception' <www.geocities.com/eganamit/puzzperc.pdf>.
- Feldman, R. (1995) 'Reliability and Justification'. *The Monist* 68: 159-74.
- Fischer, J. M. and Ravizza, M. (1998) *Responsibility and Control*. Cambridge: Cambridge University Press.
- Frankfurt, H. (1969) 'Alternate Possibilities and Moral Responsibility'. *Journal of Philosophy* 66: 820-39.
- Graff, D. (2001) 'Phenomenal Continua and the Sorites'. *Mind* 110(440): 905-935.
- Grice, H.P. (1961) 'The Causal Theory of Perception', *Proceedings of the Aristotelian Society Supplementary* 35: 121-52.
- Harman, G. (1990) 'The Intrinsic Quality of Experience'. In *Philosophical Perspectives* Vol. 4, ed. J. Tomberlin. Atascadero, CA: Ridgeview Publishing Co.
- Hawthorne, J. and Kovakovich, K. (2006) 'Disjunctivism'. *Proceedings of the Aristotelian Society Supplementary* 80: 145-183.
- Hinton, J. M.. (1967) 'Visual Experiences'. *Mind* 76: 217-27.
- Jackson, F. (1977) *Perception: A Representative Theory*. Cambridge: Cambridge University Press.
- Johnston, M. (2004) 'The Obscure Object of Hallucination'. *Philosophical Studies* 120: 113-183.
- Langsam, H. (1997) 'The Theory of Appearing Defended'. *Philosophical Studies* 87: 33-59.
- LeCatt, B. (1982) 'Censored Vision'. *Australasian Journal of Philosophy*, 60: 158-62.
- Lewis, D. (1986a) *Philosophical Papers: Vol. 2*, Oxford: Oxford University Press.
- Lewis, D. (1986b) 'Veridical Hallucination and Prosthetic Vision', in Lewis (1986a), 273-86. Reprinted from the *Australasian Journal of Philosophy* 58: 239-49. (1980).

- Lewis, D. (1986c) 'Causation', in Lewis (1986a), 159-72. Reprinted from the *Journal of Philosophy* 70: 556-67. (1973).
- Lormand, E. (2004) 'Explanatory Stopgap'. *Philosophical Review* 113: 303-357.
- Ludlow, P. (1995) 'Externalism, Self-Knowledge, and the Prevalence of Slow-Switching'. *Analysis* 55: 45-9.
- McDowell, J. (1982) 'Criteria, Defeasibility, and Knowledge'. *Proceedings of the British Academy*, 68: 455-479.
- McFetridge, I.G. (1985) 'Supervenience, Realism, and Necessity'. *Philosophical Quarterly* 35: 245-58.
- McGinn, C. (1982) *The Character of Mind*. Oxford: Oxford University Press.
- McGinn, C. (1988) 'Consciousness and Content'. *Proceedings of the British Academy* 74: 219-39.
- McGinn, C. (1989) *Mental Content*. Oxford: Blackwell Publishers.
- McGinn, C. (1999) *Knowledge and Reality*. Oxford: Clarendon Press.
- McLaughlin, B. (1984) 'Perception, Causation and Supervenience', *Midwest Studies in Philosophy*, 9: 569-91
- McLaughlin, B. (1996) 'Lewis on What Distinguishes Perception from Hallucination'. In *Perception*, ed. K. Akins. Oxford: Oxford University Press.
- Martin, M.G.M. (1997) 'The Reality of Appearances'. In *Logic, Thought, and Ontology*, ed. M. Sainsbury. Milan: FrancoAngeli.
- Martin, M.G.F. (2003) 'Particular Thoughts and Singular Thought'. In *Thought and Language*, ed. A. O'Hear. Cambridge: Cambridge University Press.
- Martin, M.G.F. (2004) 'The Limits of Self-Awareness'. *Philosophical Studies*, 120: 37-89.
- Martin, M.G.F. (2006) 'On Being Alienated'. In *Perceptual Experience*, eds. T. Gendler and J. Hawthorne. Oxford: Oxford University Press.
- Milner, A.D. & Goodale, M.A. (1995) *The Visual Brain in Action*. Oxford: Oxford University Press.
- Price, H. H. (1932) *Perception*. London: Methuen & Co. Ltd.

- Pylyshyn, Z. (2003) *Seeing and Visualizing*. Cambridge, Mass: MIT Press.
- Robinson, H. (1994) *Perception* London: Routledge.
- Russell, B. (1905) 'On Denoting'. *Mind* 14: 479-493.
- Searle, J. (1983) *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.
- Searle, J. (1991) 'Replies to My Critics'. In *John Searle and His Critics*, eds. E. LePore and R. Van Gulick. Oxford: Basil Blackwell.
- Searle, J. (2004) 'The Phenomenological Illusion'.
<<http://socrates.berkeley.edu/~jsearle/PhenomenologicalIllusion.pdf>>.
- Siegel, S. (MS.) 'Comments on David Chalmers' 'Perception and the Fall from Eden'.
<<http://www.people.fas.harvard.edu/~ssiegel/papers/edencomments.htm>>.
- Snowdon, P. F. (1980-81) "Perception, Vision and Causation". *Proceedings of the Aristotelian Society* 81: 175-192.
- Snowdon, P.F. (1990) 'The Objects of Perceptual Experience'. *Proceedings of the Aristotelian Society Supplementary* 64: 121-150.
- Soames, S. (2002) *Beyond Rigidity: The Unfinished Semantic Agenda of Naming and Necessity*. New York: Oxford University Press.
- Soteriou, M. (2000) 'The Particularity of Visual Perception'. *European Journal of Philosophy* 8: 173-189.
- Speaks, J. (forthcoming) 'Transparency, intentionalism, and the nature of perceptual content'. *Philosophy and Phenomenological Research*.
- Strawson, G. (1994) 'The Impossibility of Moral Responsibility'. *Philosophical Studies* 75: 5-24.
- Tye, M. (1982) 'A Causal Analysis of Seeing'. *Philosophy and Phenomenological Research* 42: 311-25.
- Tye, M. (2002) 'Representationalism and the Transparency of Experience'. *Nous* 36: 137-151.
- Walton, K. (1984) 'Transparent Pictures: On the Nature of Photographic realism'. *Critical Inquiry*, 11: 246-276.
- Wedgwood, R. (2002) 'The Aim of Belief'. *Philosophical Perspectives* 16: 267-297.

Wright, C. (2002) '(Anti-)Sceptics Simple and Subtle: G.E. Moore and John McDowell'.
Philosophy and Phenomenological Research 65: 330-48.