McDOWELL, GETTIER, AND THE BIPARTITE ACCOUNT OF PERCEPTUAL KNOWLEDGE

Adrian Avery Archer

A Thesis Submitted for the Degree of MPhil at the University of St. Andrews



2008

Full metadata for this item is available in the St Andrews Digital Research Repository at: https://research-repository.st-andrews.ac.uk/

Please use this identifier to cite or link to this item: <u>http://hdl.handle.net/10023/511</u>

This item is protected by original copyright

This item is licensed under a <u>Creative Commons License</u>



MCDOWELL, GETTIER, AND THE BIPARTITE ACCOUNT OF PERCEPTUAL KNOWLEDGE

A Dissertation Submitted in Partial Fulfillment of the Requirements for the degree of

MASTER OF PHILOSOPHY

By

Adrian Avery Archer School of Philosophical and Anthropological Studies

> SCOTLAND, UK August 2007

ABSTRACT

McDowell, Gettier, and the Bipartite Account of Perceptual Knowledge

Adrian Avery Archer

In his essay, "Knowledge and the Internal Revisited", John McDowell claims that "seeing that *p* constitutes false-hood excluding justification for believing that *p*." In this thesis I attempt to construct an account of perceptual knowledge that exploits McDowell's notion of false-hood excluding justification. To this end, I limn a *justified (strong) belief*, or bipartite, account of perceptual knowledge in which justification is seen as factive. On this picture, the truth requirement of the traditional tripartite account is incorporated into the justification condition for knowledge. My account of perceptual knowledge is McDowellian in spirit, but not in detail. Specifically, I part ways with McDowell in my insistence that knowledge should be seen as a composite rather than primitive concept in which belief, understood as commitment to the truth of a proposition, and justification, understood as the possession of a factive reason, both figure.

Keywords: Epistemology, Perceptual Knowledge, Philosophy of Perception, John McDowell, Selfknowledge

> Supervisors: Simon Prosser and Crispin Wright

Examining Committee: Duncan Pritchard and Martin Smith I, Adrian Avery Archer, hereby certify that this thesis, which is approximately 36,000 words in length, has been written by me, that it is the record of work carried out by me and that it has not been submitted in any previous application for a higher degree.

Date

Signature of candidate

I was admitted as a research student in September 2005 and as a candidate for the degree of MPhil in August 2006; the higher study for which this is a record was carried out in the University of St Andrews between 2005 and 2007.

Date

Signature of candidate

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of MPhil in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date

Signature of supervisor

In submitting this thesis to the University of St Andrews I understand that I am giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. I also understand that the title and abstract will be published and that a copy of the work will may be made and supplied to any bona fide library or research worker, that my thesis will be electronically accessible for personal or research use, and that the library has the right to migrate my thesis into new electronic forms as required to ensure continued access to the thesis. I have obtained any third-party copyright permissions that may be required in order to allow such access and migration.

Date

Signature of candidate

TABLE OF CONTENTS

1.	INTR	CODUCTION	1	
2.	GETTIER AND JUSTIFICATION RELIABILISM			
	2.1	Gettier and the Tripartite Account	3	
	2.2	Gettier and Justification Externalism	8	
	2.3	Gettier and Local Justification Reliabilism	12	
	2.4	Gettier and the (Anti) Anti-luck Diagnosis	16	
	2.5	Gettier and Global Justification Reliabilism	18	
	2.6	Chapter Summary	20	
3.	THE MCDOWELLIAN REPLY TO GETTIER			
	3.1	McDowellian Justification Internalism	23	
	3.2	McDowellian Content Externalism	25	
	3.3	McDowellian Singular-thought Disjunctivism	27	
	3.4	McDowellian Justificatory Dogmatism	30	
	3.5	McDowellian Internalism and Externalism Reconciled	33	
	3.6	Chapter Summary	35	
4.	SING	36		
	4.1	Infallibility and Authoritative Self-Knowledge	38	
	4.2	The First MacDonald Objection	44	
	4.3	The Second MacDonald Objection	46	
	4.4	The Third MacDonald Objection	50	
	4.5	(Re)Defining Authoritative Self-Knowledge	53	
	4.6	Chapter Summary	55	
5.	THE LOTTERY ARGUMENT FOR FACTIVE JUSTIFICATION		57	
	5.1	Factive Justification and Truth-Guaranteeing Reasons	58	
	5.2	Factive Justification and Lottery-type Beliefs	60	
	5.3	The Justification Closure Objection	62	
	5.4	The Introspective Failure Objection	64	
	5.5	The Covert Lottery Objection	67	
	5.6	Chapter Summary	73	
6.	THE BIPARTITE ACCOUNT OF PERCEPTUAL KNOWLEDGE			
	6.1	Belief and the Bipartite Account	74	
	6.2	Belief/Knowledge Disjunctivism	75	
	6.3	Strong/Weak Belief Disjunctivism	77	
	6.4	Clarifications of the Bipartite Account	80	
	6.5	Applications of the Bipartite Account	85	
7.	CON	CONCLUSION		
REFE	ERENCE	S	90	

My main point in "Knowledge and the Internal" is to protest against an interiorization of the justifications available to us for claims about the external world. The interiorization threatens to deprive us of the justificatory power of, for instance, the form "I see that" I insist that statements of such forms are proper moves in the game of giving reasons, and their truth fully vindicates entitlement to the embedded propositions. This ought to seem sheer common sense, and it would if questionable philosophy did not put it at risk.

--John McDowell, 2002, page 98.

1 INTRODUCTION

The present disquisition may be summarised in a single McDowellian thesis: 'seeing that p constitutes false-hood excluding justification for believing that p'(McDowell [2002], p. 98). What follows is an articulation of the theory of perceptual knowledge that follows from this claim. My dissertation will be McDowellian in spirit, but not in detail. Specifically, I will be co-opting some of McDowell's theoretical apparatus in defence of a bipartite conception of perceptual knowledge. That is, I see perceptual knowledge as best analysed in terms of justified *strong*¹ belief (henceforth abbreviated as *JB*). Of course, McDowell does not endorse a JB account since he, like Timothy Williamson and others, sees knowledge as a primitive. However, I believe that the central intuitions motivating the primitivist account of knowledge may be captured within my own JB account of perceptual knowledge, while avoiding certain problems inherent in the primitivist view. In many respects my position resembles the more traditional tripartite or *justified true belief* (henceforth abbreviated as *JTB*) conception of knowledge. However, it is central to my view that justification, at least with

¹ The significance of the qualification 'strong' is explained in §6.3

regards to our perceptual beliefs, is factive. (Thus, on my proposed account, the 'T' of the traditional JTB account is incorporated into the 'J'.) The central contention of my dissertation is that perceptual knowledge may be analysed in terms of justified (strong) belief, where justification is seen as a factive notion.

The emphasis on perceptual knowledge is crucial since I will not be concerned with knowledge *simpliciter*. Thus, I remain open to the possibility that other forms of (non-perceptual) knowledge do not conform to the JB analysis limned in the present study. By perceptual knowledge I simply mean knowledge that is based on perceptual experience in the following sense: if one were asked, why one believed that p, one's answer would implicate a perceptual experience. That is to say, if the answer one would give to the question, 'why does one believe that p?', is that one sees that p, hears that p, smells that p or tastes that p, then on my view one's belief that p is a perceptual belief. Moreover, if one's perceptual belief that p is justified, then on my account one has perceptual knowledge that p. For simplicity, my discussion will be limited to the case of vision, but presumably what follows may be generalised with respect to the other sensory modalities. Moreover, I take perceptual knowledge to be propositional in content. Hence, one may see, for example, *that there is a vase on the table*. In what follows, whenever I speak of a subject seeing that p, it is propositions of this form that I have in mind.

Thematically, my dissertation may be divided into two halves. The first half, §2-§4, represents an exposition and defence of key features of McDowell's theory of knowledge. I will be using a pair of perceptual Gettier cases as a foil for bringing into sharp relief four McDowellian theses: (1) justification internalism, (2) content externalism, (3) singular-thought disjunctivism and (4) justificatory dogmatism. The second half, §5 and §6, represents my appropriation of these four theses in defence of a bipartite account of perceptual knowledge.²

² I will like to thank Katherin Hawley, Duncan Pritchard, Simon Prosser, Martin Smith and Crispin Wright for their valuable feedback at various stages of this dissertation.

2 GETTIER AND JUSTIFICATION RELIABILISM

2.1 Gettier and the Tripartite Account

If there were anything like a classical or traditional definition of knowledge, it would be something along the following lines:

JTB	For any subject <i>S</i> , <i>S</i> knows that p <i>iff</i> ³ :	(i)	<i>p</i> is true
		(ii)	<i>S</i> believes that <i>p</i>
		(iii)	S's belief that p is justified

JTB takes (i), (ii), and (iii) to be individually necessary and jointly sufficient for *s*'s knowing that *p*. That knowledge is composed of these three features was an assumption taken for granted by philosophers dating at least as far back as Plato.⁴ However, this fact should not prompt us to dismiss JTB since it also appears to receive the endorsement of common sense. First, it seems counterintuitive to talk about knowing that *p* if *p* is false. Second, one cannot know that *p* if *p* is not among one's beliefs. Third, knowing that *p* requires more than simply believing that *p*. It also requires one to have some *reason* for, *evidence* in favour of, or *justification* for believing that *p*. (For the time being, I will treat all these notions as synonymous.) Significantly, the third requirement of JTB has traditionally been taken to express a J-internalist intuition—namely, that one's reasons should be constituted by factors that are subjectively available to one.⁵ Each of these theses appears to be in harmony with our everyday usage and understanding of the word 'know'. Moreover, one would be hard pressed to find some additional requirement necessary for our quotidian understanding of knowledge.

³ The '*iff*' is to be understood as indicating necessary coextenstion, or equivalence.

⁴ See: Plato [1961]: 97e-98a.

⁵ The third requirement of JTB can also be spelt out in J-externalist terms. However, J-internalist accounts of justification tend to do a better job at preserving the common sense understanding of justification as being in a position to provide reasons for one's beliefs.

Unfortunately, Gettier-type counterexamples have raised serious questions about the adequacy of JTB.⁶ For example, consider the case of veridical hallucination below.

2.1.1 The Gettier Case of Veridical Hallucination

Suppose *s* has strong perceptual evidence for, and comes to believe, the empirical proposition:

(a) there is a vase on the table.

S's evidence for (a) is that she is having a visual experience of there being a vase on the table. But unknown to S, her visual experience of there being a vase on the table is actually a druginduced hallucination. Now, it so happens that there is in fact a vase on the table, though the real vase is being obscured from S's visual field by a stack of books. The following all seem to be true of the aforementioned case:

- (i) (a) is true
- (ii) *S* believes that (a)
- (iii) *S*'s belief that (a) is justified

Thus, all three conditions of JTB have been met. However, intuitively we would be reluctant to say that the subject in the above example knows that (a). (I will refer to the Gettier case of veridical hallucination just described with the abbreviation $GTTR_{VH}$.) There are a number of explanations that could be given for the intuition that the $GTTR_{VH}$ subject lacks knowledge that (a). For example, we may say that *S*'s reasons fail to constitute *good* reasons since the evidential basis for her belief that (a) is a hallucination rather than an actual vase. Put differently, we may say that *S* does not have the reasons she takes herself to have for believing that (a). Or, we may point out that there is no objective connection between *S*'s justification for believing that (a) and the fact itself. For example, we may say that *S*'s belief

⁶ See Edmund Gettier [1963].

was not caused by the relevant fact. Or we may say that the fact that *S*'s belief that (a) happens to be true is (given the circumstances) merely a matter of luck; she would have believed (a) even if it were true that not-(a).⁷ The first explanation is likely to be given by a J-internalist, while the second would probably be offered by a J-externalist. The third explanation, which attributes *S*'s knowledge-failure to luck, is the most neutral of the three. Whether J-internalist or J-externalist, most philosophers agree that the luck involved has something to do with the intuition that the Gettier subject does not have knowledge.

My proposed bipartite account presupposes an alternative framework to the standard anti-luck diagnosis of why the Gettier subject lacks perceptual knowledge. I take the anti-luck diagnosis to be inadequate for one of two possible reasons. On the one hand, there seems to be genuine Gettier cases in which it is not a matter of luck that the subject's belief is true. On the other hand, there appear to be cases in which we would intuitively impute perceptual knowledge to a subject even though the fact that her belief is true is partly due to luck.⁸ Pace the anti-luck diagnosis, I argue that the Gettier subject lacks perceptual knowledge that (a), there is a vase on the table, because she does not have the (truthguaranteeing) reason, she takes herself to have, for believing that (a). Let us call S's reason for believing that there is vase on the table {R}. In the aforementioned case, {R} = seeing that there is a vase on the table. However, since S does not see that there is a vase on the table, she does not have $\{R\}$ – i.e., the (truth-guaranteeing) reason she takes herself to have. It is still too early to fully appreciate this point, but I will return to it in §6. My proposed account attempts to preserve the J-internalist intuition that one must have a 'good reason' (in which a reason is seen as something one has subjective access to) for believing that p in order to be justified in believing that *p*. However, I also attempt to preserve the J-externalist intuition that there must be an objective connection between justification and truth.

There are three strategies available for responding to Gettier presently represented in the literature:

⁷ Gettier has only shown that the internalist requirement is not a sufficient condition for knowledge. It has not shown that is not a necessary condition for knowledge.

⁸ Which of these two rationales for why the anti-luck diagnosis is inadequate we find preferable depends on our understanding of epistemic luck. I will not attempt to settle which of the rationales is best, since both lend equal support to my conclusion.

- **(REPLY2)**: Accept that Gettier is correct and attempt to come up with a new conceptual analysis of knowledge.
- **{REPLY3}**: Deny that Gettier cases (like GTTR_{VH}) are instances of justified belief.

In its most extreme form, {REPLY1} amounts to the claim that practically all true belief, including lucky guesses and Gettier cases, constitute knowledge. This is the position of epistemic minimalists, such as Crispin Sartwell, who dispense with the concept of justification altogether. We may describe such approaches as offering a TB account of knowledge. TB accounts of knowledge are regarded as highly implausible by most analytic philosophers since they require a substantial revision of our putative conception of knowledge. Since TB accounts involve the outright rejection of the intuitions underlying the justification requirement, it is not a position we need to take seriously given our present goals—namely, a defence of an intuitively plausible picture of justification.

{REPLY2} is arguably the most popular among contemporary epistemologists and has spawned at least three separate types of replies to Gettier. First, there are the TB+G accounts of knowledge, according to which knowledge is composed of true belief plus some codicillary anti-Gettier condition. TB+G approaches include Alvin Goldman's [1967] 'causal condition' and Keith Lehrer and Thomas Paxson's [1969] 'defeasibility condition' analyses. What these approaches all share is that they posit some extra anti-Gettier requirement that will yield a set of necessary and jointly sufficient conditions for knowledge. Second, there are the JTB+G accounts of knowledge. One example of the JTB+G approach is the "no false lemmas" analysis which rests on the observation that in most extant Gettier examples, the justified true belief came about as the result of entailment from justified false beliefs (e.g., *S*'s belief that she sees a vase). Thus, according to the 'no false lemmas' approach, we must add a forth condition to JTB stipulating that the justified true belief should depend on no false premises. Such a view has been suggested, though perhaps not explicitly endorsed, in Meyers and Stern [1973], Armstrong [1973, p. 152], and Harman [1973, p. 47]. I take the nofalse-lemmas condition for knowledge to be highly plausible and will be appealing to a version of it in my proposed JB account. Third, there are approaches that reject the attempt to construct knowledge out of set of composite properties altogether. These so-called 'knowledge-first' approaches view knowledge as conceptually basic, irreducible, or primitive. One example of such an approach is Timothy Williamson's [2000] conception of knowledge as a factive (mental) state. This is also the position favoured by McDowell. However, I view McDowell's rejection of composite accounts to be premature, particularly since the tools made available by his own approach allows us to preserve just such an account.⁹

(REPLY3) attempts to preserve JTB by modifying the justification condition in a way that prevents a subject's beliefs from being Gettiered. The most popular, though still controversial, instantiation of (REPLY3) is J-reliabilism which redefines justification in terms of reliability. Since I take J-reliabilism to be the main competitor to my own proposed account I will have a great deal to say about it later. Significantly, the JTB and TB+G forms of reliabilism are extensionally equivalent—to wit, whatever counts as knowledge on the former also counts as knowledge on the latter, and vice versa. However, there is still an important difference between the two. Consider some belief θ which is formed via a reliable process but is nevertheless false. Both the JTB and TB+G accounts agree that, given the factivity of knowledge, the proposition implicated in θ is not known. However, JTB reliabilism also implies that θ is justified, even though it fails to qualify as knowledge. TB+G reliabilism is silent on the question of justification. (Many would take this to be a strength of JTB reliabilism over TB+G reliabilism since they regard justification as an interesting concept in its own right, independent of its relationship to knowledge.)

My own JB account of knowledge represents a hybrid of {REPLY2} and {REPLY3}. Like advocates of {REPLY2} I concede that Gettier is right about the insufficiency of the

⁹ In §6.1 I present a number of difficulties facing the primitivist account, to which my own bipartite account is immune. It should also be noted that most philosophers believe that the introduction of primitives into a philosophical theory is best left as a last resort. However, this is not an issue I intend to take up in this paper.

traditional JTB account of knowledge. I therefore offer an alternative bipartite account. However, like {REPLY3}, my account maintains that Gettier cases (at least with respect to our perceptual beliefs) are not cases of justified belief. The main lesson I draw from McDowell's theory of knowledge is that truth cannot be taken to be external to a subject's epistemic standing if we are to arrive at a satisfactory account of perceptual knowledge. The upshot of this idea for my JB account is that justification must entail truth. However, the notion of truth-entailing or factive justification may initially seem unnecessarily strong. Moreover, it has the (apparently) counter-intuitive result that a justified belief cannot be false. Thus, my proposed JB account of perceptual knowledge may seem unmotivated. This is one of several objections I address in §6. Presently, I will be examining a competing theory of justification in terms of reliability.

2.2 Gettier and Justification Externalism

J-externalists blame the failure of JTB on the J-internalist notion of justification defined purely in terms of subjectively available reasons for one's beliefs. According to this line of thought, subjectively available reasons are denizens of the 'inner' world (or the mind) and are therefore independent of, and insulated from, the 'external' world. The upshot of this picture, as Williamson [2000] notes, is that the mind and the external world are characterised as two independent variables [p. 5]. Having a justified belief is simply a function of the 'inner world' (or mind) while the truth of a given empirical proposition remains a function of the 'external world'. However, J-internalism fails to provide an account of how there could be a non-accidental connection between the two. On the J-internalist picture, any similarity between the content of the inner world and that of the external world—between belief and truth—seems to amount to happenstance or luck, or so the argument goes.

The key to escaping the unfavourable implications of GTTR_{VH}, according to Jexternalists, is to define knowledge in such a way that a subject's perceptual beliefs are connected, not to internally fixed reasons but to the external world itself. The most widely discussed (if not accepted) version of J-externalism is J-reliabilism¹⁰, which asserts that for a belief that p to be justified it must reliably indicate the fact that p, or be reliably connected to the fact that p, or be produced by a reliable method or process. For simplicity, I will limit myself to the process version of J-reliabilism:

Process J-Reliabilism:

For any subject *S*, *S*'s belief that p is justified *iff* it was formed via a reliable process (i.e., a process that tends to produce true beliefs).¹¹

There are two dimensions to the notion of reliability. On the one hand, we may say that some process {M} is reliable so long as it tends to produce true beliefs with respect to some particular proposition, p. This would be a description of what McGinn and Goldman refer to as *local* reliability. Local reliability contrasts with what has come to be called *global* reliability. On this view, a process, {M}, is globally reliable when it tends to produce true beliefs for a range of p-like propositions.¹² Thus, global reliability requires that {M} be reliable with respect to not only the particular proposition p but also with respect to a larger range of propositions. As Goldman explains it, the notion of global reliability is a 'statistical or dispositional property of a belief-forming process or type' [p. 49]. A type of belief-forming process is globally reliable *iff* it tends to produce a high ratio of true beliefs for a range of uses of that process. For the time being I will only be concerned with local reliability. (I will return to the question of global reliability in §2.4.) We may reword the above definition of process J-reliabilism in order to capture the local nature of the reliability in question:

¹⁰ Reliabilism is typically taken by McDowell as the paradigm instantiation of epistemic externalism. I follow him in this regard.

¹¹ See Goldman [1979] for the *locus classicus* of this view. (J-Rel_{PR}) and other forms of J-reliabilism should be distinguished from the TB+G reliabilist accounts which do not represent authentic defences of JTB. TB+G reliabilism omits all talk of evidence or justification and instead conceives of knowledge as true belief that satisfies a causal or reliability condition. On this view reliability is what you add to true belief *instead* of evidence or justification, in order to arrive at knowledge. Examples of this version of reliabilism include Goldman's [1967] 'causal condition' approach and Robert Nozick's [1981] truth tracking account. TB+G accounts are versions of K-externalism, since they are concerned with knowledge rather than justification. J-reliabilism, by contrast, views reliability not as replacing the role of justification in an account of knowledge but rather as that which *makes* a belief justified. (J-Rel_{PR}) therefore represents a bona fide defence of JTB. ¹² See Goldman [1986] p. 45.

⁹

Local Process J-Reliabilism:

For any agent *S*, *S*'s belief that *p* is justified *iff* it was formed via a process that is locally reliable. (i.e., a process that tends to produce true beliefs with respect to a particular proposition, *p*)

Let us assume, for the sake of argument, that knowledge is true belief acquired by a process, {M}, which is locally reliable. It remains an open question as to exactly what {M} has to be reliable for. Here are a few possibilities¹³:

- (R1) For producing a true belief with regards to a particular proposition *p* in precisely the situation in question.
- (R2) For producing a true belief with regards to a particular proposition *p* in situations just like the one in question.
- (R3) For producing a true belief with regards to a particular proposition *p* in all situations worth considering.
- (R4) For producing a true belief with regards to a particular proposition *p* in all possible situations.

Notice that (R1)-(R4) all relate to a particular proposition p, and are therefore all versions of local reliability.¹⁴ However, each subsequent statement increases the range of possible situations in which a given process, {M}, must work in order for the subject to count as knowing the particular proposition p. As definitions of what it means for a process to be reliable, (R1) and (R2) strike me as trivial. Both, more or less, amount to the claim that {M} is reliable when {M} is reliable (though whether we regard (R2) as trivial depends on how

¹³ See Craig [1990, p. 54] for a similar list.

¹⁴ This should not be confused with the increasing strength of the reliability requirement embodied in the shift from local reliability (i.e., regarding a particular proposition p) to global reliability (i.e., regarding all relevantly p-like propositions). A similar list can be constructed on behalf of global reliability by adjusting the notion of relevant likeness to include a progressively wider range of relevant p-like propositions.

narrowly we understand the 'just like' in 'situations *just like* the one in question').¹⁵ This is not to suggest that (R1) and (R2) are trivial vis-à-vis a definition of justification. In fact, as a definition of justification (R1) is no less demanding than (R4) since it implies one is only justified if one's belief that *p*, in a particular instance, is factive. The most widely discussed version of reliability is (R3). The locution, 'situations worth considering' that features in (R3) raises important questions that bear on the debate regarding relevant alternatives and how we should go about settling which situations are close enough to the one in question. For example, are we to think of relevance as designating a similarity class (i.e., situations that resemble each other in certain fundamental respects) or as an interest claim (i.e., situations that appeal to us because of the context of discussion or particular goal we have in mind).¹⁶ These are questions that must be settled if we are to arrive at a workable understanding of local J-reliabilism. However, they do no bear directly on our present discussion and so we may table them for now.

The application of local J-reliabilism to the Gettier case of veridical hallucination described earlier is quite straightforward. Hallucinating a vase is not a reliable method for arriving at true beliefs about vases. In the case of the subject who comes to believe (a), *there is a vase on the table*, based on a hallucination of a vase, the belief that (a), though true, is not reliably formed. Thus, according to local J-reliabilism, the GTTR_{VH} subject does not have a justified belief. However, below I will attempt to construct a Gettier case against local J-reliabilism. Since (R3) is the most widely accepted version of local J-reliabilism, my proposed Gettier case will take it as its target. I will then consider a strengthened version of local J-reliabilism, one that appeals to Nozick-style counterfactual conditionals.

My motivation for arguing against a Nozick-style version of J-reliabilism is two-fold. First, Nozick-style J-reliabilism implicates factive justification and therefore represents a Jexternalist competitor to my proposed J-internalist friendly version of factive justification. While my bipartite account of perceptual knowledge also implicates factive reasons, it contrasts with Nozick-style justification in that it sees justification as rooted in a subject's

¹⁵ This is not to suggest that (R1) and (R2) are trivial vis-à-vis a definition of justification. In fact, as a definition of justification (R1) is no less demanding than (R4) since it implies one is only justified if one's belief that p, in a particular instance, is factive.

¹⁶ Which situations count as 'similar' also raises issues relating to what features of a given situation should count as salient vis-à-vis a subject's justification.

possession of reasons. In this regard, it differs from the extreme J-externalist accounts, such as J-reliabilism, that omit all talk of reasons.¹⁷ Second, by showing that Nozick-style Jreliabilism is insufficient for responding to Gettier, I hope to add support to my claim that luck-based diagnoses of Gettier are inadequate. Since it is not simply a matter of luck that the subject with Nozick-style justification has true beliefs (a claim which I qualify below) the fact that such a subject may be Gettiered suggests that the presence of luck is not the sole explanation of what may go wrong in Gettier-type cases.

2.3 Gettier and Local Justification Reliabilism

Consider the following perceptual Gettier case. Bob enters a classroom and sees what he takes to be a vase in the middle of the floor. Based on this perceptual experience, Bob forms the following belief:

(a*) There is a vase in the room.

Bob believes that *there is a vase in the room* because he thinks that he sees that *there is a vase in the room*. However, what Bob takes to be a vase is actually a computer generated hologram. Now, there is in fact a real vase in the room, but the actual vase is hidden behind an opaque black curtain at the back of the room. Furthermore, the computer that generates the hologram is equipped with special sensors that detect the presence of the vase behind the curtain. This then prompts the computer to project the hologram of the vase in the centre of the room. In short, the hologram in the centre of the room is causally connected to the vase hidden behind the curtain at the back of the classroom. Moreover, the computer's programming is perfectly reliable, so that once there is a vase in the room. In sum, as far as the computer program is concerned, the presence of the vase in the centre of the room is a necessary and sufficient condition for the hologram of the vase in the centre of the room. Unfortunately (or perhaps fortunately), Bob lacks any of this background information,

¹⁷ Thus, my discussion of J-reliabilism may be seen as a foil for bringing into sharp relief the J-internalist friendly nature of my proposed bipartite account.

and forms his belief that there is a vase in the room based solely on the hologram in the centre of the room. All of the following seem true in the above case:

- (i*) (a*) is true
- (ii*) Bob believes (a*) is true
- (iii*) Bob's belief that (a*) is formed via a reliable process

Ex hypothesi, (iii*) is true since the computer is programmed to only project the hologram of a vase when there is an actual vase in the room. Thus, Bob's belief that there is a vase in the room is locally reliable since the process by which the belief was formed would, given the computer's programming, tend to produce true beliefs with respect to the particular proposition (a*). However, I believe the above scenario represents a bona fide Gettier case since, though Bob's belief is justified (i.e., reliably formed) and true, we would not want to say that he has knowledge.¹⁸ I will refer to the Gettier case against local reliability just described with the abbreviation *GTTRLR*.

It seems clear that local J-reliabilism, when understood in terms of (R1) is insufficient for resisting GTTRLR. According to (R1), a process, {M}, is reliable so long as it is reliable with regards to the situation in question. The computer program described in GTTRLR clearly meets this stipulation. (R2), by contrast, is concerned not only with the situation in question, but also with situations 'just like it'. Thus, (R2) seems to introduce the idea of some sort of unspecified counterfactual condition for a process being reliable. This notion takes on a much more robust form in (R3) where a range of possible situations seems to be implicated.

2.3.2 Gettier and Nozick-Style Justification

We may further unpack the idea of counterfactuals, hinted at in (R2) and suggested by (R3), in terms of Nozick's [1981: 179] subjunctive conditionals:

¹⁸ I believe Bob does not have perceptual knowledge that there is a vase in the classroom because he does not see the vase. On this score, I am in agreement with Paul Snowdon's observation that seeing that p puts one in a position to have demonstrative thoughts about p. If Bob attempted to have demonstrative thoughts about the vase he would fail to do so since he would point to the hologram rather than the real vase. As such, Bob does not see the vase.

- (N1) If *p* were not true, then *S* would not believe that *p*
- (N2) If p were true, then S would believe that p^{19}

Together, (N1) and (N2) allow a subject's belief to track the truth of a given proposition, which has earned Nozick's version of reliabilism the title: the truth-tracking account. Significantly, Nozick's truth-tracking account is actually a type of TB+G reliabilism and is therefore not a genuine version of J-reliabilism. However, we can, for the time being, ignore this fact by thinking of (N1) and (N2) as the necessary and sufficient conditions for justifiably believing that p. We may say that (N1) and (N2) are true, as counterfactual statements, *iff* in possible worlds near to the actual world, if *p* is false, *S* does not believe that *p*, and if *p* is true *S* believes that p^{20} Nozick recommends that we assess (N1) and (N2) by reference to what is the case in all nearby possible worlds and not just the actual world. Roughly, a world may be described as 'nearby' if it is only slightly different from the actual world and 'distant' if it is radically different. Given the above account of nearby worlds, the following reply to GTTR_{LR} seems available to Nozick. Let us suppose that the set of nearby possible worlds include ones in which the computer crashes or malfunctions. In such nearby worlds, although there is a vase in the room, there is no hologram of a vase. Since in such a world the subject would not believe that (a^*) although (a^*) were true, (N1) has not been satisfied.

There are a number of reasons why the above reply to GTTR_{LR} is unsatisfactory. First, there does not seem to be any principled reason for holding that the set of nearby worlds include ones in which the computer crashes or malfunctions. For example, suppose the computer in question is of a particularly high quality (we may suppose that it is a Mac rather than a PC). Would a world where the computer crashes or malfunctions still count as

¹⁹ Counter examples such as the Grandma case has prompted Nozick [1981, p. 179] to revise (N1) and (N2), limiting them to the same method (i.e., vision):

⁽N1*) If p were true, S (using M) would believe that p

⁽N2*) If p were not true, then S (using M) would not believe that p

 $⁽N1^*)$ and $(N2^*)$ restricts us, not only to a particular proposition p, but also to a particular method M for arriving at p. Since I am combining Nozick's conditionals with process reliabilism, the restriction to a single method is already built into the present analysis.

²⁰ I appeal to the argot of possible worlds only as a linguistic convenience and nothing more. For further clarification of the semantics of subjunctive conditionals see Stalnaker [1968]; Lewis [1973] and Nozick [1981] note 8.

nearby? This question points towards a certain capriciousness in the determination of nearby worlds. Consider the following often cited example:

- (A) If Nixon had pressed the red button there would have been a nuclear war.
- (B) If Nixon had pressed the red button the connecting wire would have broken.

On at least one interpretation, the world in which (B) obtains is more similar to the actual world in that it, like the actual world, has not undergone a nuclear war. However, the possibility of the wire breaking does not seem salient given the context suggested by the content of the two sentences. Thus, intuitively, (A) seems like the acceptable counterfactual, while (B) seems both arbitrary and unwarrantedly optimistic. The question that confronts the nearby worlds type reply to GTTR_{LR} is why should we consider the computer malfunctioning salient vis-à-vis the determination of nearby worlds, particularly once we have already stipulated that the computer in question is not one prone to malfunctions. Of course it would best serve Nozick's purposes to do so, but such a motivation would seem, at best, ad hoc.

But even more crucially, Nozick seems to have good reason for limiting the set of nearby worlds to ones in which the computer is working properly. For example, while not common, people do occasionally go blind due to damage or disease. Now, if the set of nearby worlds is taken to include ones in which a perfectly healthy subject suddenly goes blind, then vision would no longer count as a reliable method. But this is simply implausible. Now conceivably, there could be a computer which is no less reliable in its operation, than a human sense organ. Thus, if worlds in which a healthy subject goes blind are excluded from the set of nearby worlds, then the same should hold for worlds in which a perfectly functioning computer malfunctions.

The point of the above example, then, is that any account which included worlds in which the computer malfunctions as nearby seems to prove too much. This is a fact that Nozick himself seems to be sensitive to, as evinced in his treatment of the Jesse James case.²¹ Here is the Jesse James example, as described by Craig [1990]:

²¹ See: Nozick [1981] p. 193.

Jesse James, the reader will recall, is riding away from the scene of the crime with his scarf tied round his face just below the eyes in the approved manner. The mask slips, and a bystander, who has studied the 'wanted' posters, recognises him. The bystander now knows, surely, that it was James who robbed the bank. But Nozick has a problem: there is a possible world, and a 'close' one, in which James' mask didn't slip, or didn't slip until he was already past the bystander; and in that world the bystander wouldn't believe that James robbed the bank, although it would still be true that he did. So Nozick's condition [N2] is not satisfied, and he is threatened with having to say that the bystander doesn't know that it was James, even though the mask did slip. So his analysis looks like ruling out something which is as good a case of knowledge as one could wish for [p. 22].

Nozick is already equipped with a reply to the 'Jesse James' counterexample, akin to that employed in the Grandma case (see footnote 19 above). In brief, he may simply argue that in worlds in which the mask did not slip, the method employed by the bystander would be different. Thus, given the version of Nozick's counterfactuals revised to include the subject's method, worlds in which the mask did not slip would not be included in the relevant nearby possible worlds. However, to the extant that this reply is effective in preserving the bystander's knowledge in the Jesse James case, it is also effective at preserving Bob's 'knowledge' that there is a vase in the room. In GTTRLR, the computer constitutes part of the method by which Bob arrives at his belief that there is a vase in the room. Thus, by the lights of (N2*), all worlds in which the computer has crashed, malfunctioned or missing altogether, a different method is being used and such worlds would *eo ipso* fail to count as nearby.

2.4 Gettier and the (Anti) Anti-Luck Diagnosis

The traditional diagnosis of what goes wrong in Gettier-type scenarios is in terms of luck. While I believe there is much to be said for the standard anti-luck diagnosis, I do not think it is the sole explanation of what can go wrong in Gettier-type scenarios. I believe this claim is supported by the fact that a subject whose belief was formed by a reliable process can still be Gettiered. If a subject's belief is reliably formed, then prima facie, it cannot simply be a matter of luck that it is true. However, it may be objected that there is still a species of luck involved in Bob's belief that *there is a vase in the room* since it is only a matter of luck that the hologram he sees is reliably linked to a real vase via a computer. Thus, we may say that Bob is lucky that his belief is reliable. Moreover, one may insist that it is this higher order luck (if you wish to put it like that) that is responsible for Bob's belief being Gettiered.

While tempting, I believe the above defence of the anti-luck diagnosis is deficient. Even if we grant that it is matter of luck that Bob's belief is reliably formed, this cannot be the explanation of why his fails to constitute perceptual knowledge. This is because there are cases implicating such "higher order" luck which still seem to constitute legitimate cases of perceptual knowledge. For example, consider a subject, Alan, who is out on a leisurely drive in an unfamiliar countryside. He comes to a fork in the road, which we may describe as {ROAD1} and {ROAD2} respectively. Unknown to Alan, both roads terminate at a different aviary. The aviary at the end of {ROAD1} contains only Robbins and Woodpeckers. The aviary at the end of {ROAD2} contains only Robbins and Sparrows. Alan cannot decide which of the two roads to take, and so he decides to flip a coin. {ROAD1} wins the coin toss.

Now suppose that Alan is able to reliably distinguish between Robbins and Woodpeckers, so that he never confuses the two. However, Alan cannot distinguish between Robbins and Sparrows and is prone to identify both species as Robbins. We may sum up this idea by saying that Alan's ability to recognise Robbins by sight is indexed to environments containing only Robbins and Woodpeckers but not to environments containing Robbins and Sparrows. Since Alan took {ROAD1} he ultimately visits the aviary with only Robbins and Woodpeckers. Alan sees a Robin and pointing to it says, 'there goes a Robin!' Since Alan can reliably distinguish between Robbins and Woodpeckers, intuitively we would say that Alan knows that the bird is a Robin. However, it is only a matter of luck that Alan finds himself in an environment to which his ability to recognise Robbins is indexed. After all, if the coin-flip had gone the other way, Alan would have gone down {ROAD2} and would have found himself in an environment in which he was not able to reliably identify Robbins by sight. Thus, we may say that it is only a matter of luck that Alan's belief is reliably formed. Nevertheless, the fact that Alan's belief that the bird before him is a Robin is lucky in this respect does not undermine our intuition that he has perceptual knowledge. I

conclude that in the Alan case, just as in the Bob example, it is only a matter of luck that the subject's belief is reliably formed. However, while intuitively the former has perceptual knowledge, the latter does not. Thus, the presence of luck cannot account for the epistemic difference between the two cases.²²

2.5 Gettier and Global Justification Reliabilism

The stronger notion of global reliability offers us a way to respond to GTTRLR. Recall, global reliability requires that the process by which *S* forms her belief that *p* be reliable with respect a range of *p*-like propositions. However, since the subject in the GTTRLR is prone to mistake holograms of vases for real vases, then the subject's reliability does not extend to other cases in which a hologram is not causally connected to an actual vase. Thus, from a global reliabilist perspective, the subject is not reliable. Thus, the stronger notion of global reliability succeeds where the weaker notion of local reliability fails. Unfortunately, global reliability's strength (pun intended) may prove to be its greatest weakness, since it is doubtful that we consider someone justified in believing that *p*, only if they are reliable regarding a range of *p*-like propositions.

More often than not, any process, {M}, that is reliable with regards to a particular proposition p, would also be reliable with regards to a number of p-like propositions as well. Even so, it is far from established that this is a conceptual requirement of what it means for some process, {M}, to be reliable. In other words, it may be that while the extension of our concept of reliability is typically global in scope, the intention is nevertheless only local in nature. Furthermore, it is only the intention of the concept that concerns us. The question, then, is this: is S' knowing that p contingent on her being reliable regarding a wider class of p-like propositions? There appear to be a least some cases in which there could be knowledge that p when the relevant class of 'p-like' truths doesn't reach beyond p itself, assessed at other times. This typically occurs when the subject stands in a unique

 $^{^{22}}$ The conclusion that the anti-luck diagnosis is not the sole explanation of what goes wrong in Gettier-type scenarios is one of the motivations for my proposed bipartite account. Unfortunately, space will not allow me to develop this idea at length here, though I will return to this issue again briefly in §6.

relationship to the fact in question, or may be presumed to be uniquely familiar with it. Edward Craig [1990] advances an argument to this effect:

An infant may know what its name is without knowing what anyone else's name is, and without knowing its age, and it may know its forename without knowing its surname. Being confident that someone knows their name doesn't help us to judge that they are valuable informants on many other questions, if any. If we try hard enough we will probably be able to imagine circumstances in which we might reasonably trust a potential informant on the question whether there was a table in the next room, but not on the question whether there were chairs there, or anything else about the furniture. (In fact one doesn't even have to try very hard: suppose that he has never been in the room, but has just seen a table being carried through the door.) So if we are looking for strictly necessary conditions...we shall plump for localism and confine ourselves to the individual proposition at issue. [p. 57]

However, more needs to be said if we are to find Craig's examples persuasive. For example, a defender of global reliability may counter as follows: The above argument seems to rest on the mistaken assumption that for a process, {M}, to be reliable regarding p, it must have actually been reliably employed in relation to a number of p-like propositions. Craig argues that the fact that we would count an infant as knowing its own name although it does not know anyone else's shows that local reliability is enough. But the fact that an infant may know its name and no one else's does not mean that the process, {M}, by which it acquired this knowledge is only locally reliable. It may very well be the case that {M} would also allow the baby to learn other names, but the infant simply hasn't had the opportunity (as of yet) to implement the method towards this end. Recall that according to Goldman, reliability is merely a dispositional feature of {M}. The question, then, is not whether {M} has been reliable with regards to a number of p-like propositions, but rather *would* it be. The same goes for Craig's example of the person who sees the chair (and only the chair) being taken into the room. Let us assume that the method by which Craig's subject comes to know that there is a chair in the room is something like the following:

$\{M^*\}$ seeing α being taken into the room (where α stands for some object or other)

Now, it so happens that the Craig subject has only been able to implement {M*} once, and so only has knowledge of the single proposition, *there is a chair in the room*. However, presumably, if the subject were to see a table or a television being taken into the room, he would also, using {M*} come to know that there was a table or television in the room. The fact that the subject has not had the opportunity to implement this method (as of yet) is beside the point.

Although Craig's objection is ultimately unsuccessful in impugning global reliabilism, I believe the problem lies in Craig's example rather than his strategy. If we could construct an example in which a subject intuitively has knowledge, despite the fact that the process by which she acquires her beliefs lacks global reliability, then perhaps we may succeed where Craig fails. I believe such an example is available. Consider a Royal Air Force officer who is using a second-rate radar to identify enemy airplanes. The radar system is designed to bleep only when there is a B-52 within a hundred yards of the military base. However, the radar cannot detect any other airplanes besides B-52s. Moreover, when it detects a B-52, it cannot tell exactly how near or far it is, the bomber's velocity or trajectory, or any other subsidiary information apart from the fact that the bomber is within a hundred yards of the base. In short, the radar is reliable with regards to the proposition: 'there is a B-52 within a hundred yards of the military base', but not with regards any other propositions. Thus described, we may say that the radar lacks global reliability. However, despite the radar's lack of global reliability, we would still say that when the radar bleeps, the Royal Air Force officer knows that there is a B-52 within a hundred yards of the base. This suggests that a subject may arrive at knowledge even though the process by which she acquires her beliefs lacks global reliability. Global reliability therefore fails to constitute a necessary condition for knowledge.

2.6 Chapter Summary

The definitive challenge to JTB are Gettier counterexamples which seem to present cases of justified true belief that fail to constitute knowledge. According to the J-externalist, the key to resisting Gettier is to posit an objective connection between justification and truth. The standard means by which such a connection is achieved is by insisting that a belief is justified only if it is formed via a reliable process. On this view, the Gettier subject lacks reliability and therefore is not justified. However, I have argued that local reliability is not sufficient for responding to all Gettier cases. I believe this fact casts some doubt on the standard diagnosis of what goes wrong in Gettier cases in terms of luck. Global reliability appears to succeed where its local counterpart fails. However, while global reliability seems sufficient for resisting Gettier, it does not appear to be necessary for knowledge. We therefore find ourselves confronted with an all too familiar epistemological dilemma, how to find a justification condition that is both sufficient for resisting Gettier, and necessary for knowledge. I believe McDowell's theory of justification offers the tools for constructing just such an account, and it is to his approach that we now turn.

3 THE MCDOWELLIAN REPLY TO GETTIER

John McDowell's account of perceptual knowledge is arguably one of the most innovative presently available. However, one risk that often accompanies being innovative is that it is sometimes difficult to see how one's ideas relate to more traditional lines of thought. In this chapter I attempt go some way towards remedying this problem. I will be attempting to make explicit, motivate and clarify, what I take to be four central pillars of McDowell's epistemological project:

- (M1) *McDowellian J-internalism*: For all subjects S_1 and S_2 and worlds W_1 and W_2 , if S_1 in W_1 and S_2 in W_2 are identical in terms of how things are with them subjectively, then S_1 and S_2 are identical in all respects relevant to the justification of their beliefs.
- (M2) *McDowellian C-externalism*: There are subjects S_1 and S_2 and worlds W_1 and W_2 , such that S_1 in W_1 and S_2 in W_2 have the same intrinsic properties but differ in the content of their thoughts.
- (M3) *McDowellian S-disjunctivism*: For any subject *S*, if it seems to *S* as if she is having the singular thought, α is φ , either the external object, α , figures in the content of her thought or she is suffering an illusion of thought.
- (M4) *McDowellian J-dogmatism*²³: For any subject *S*, if *S* is justified in believing some perceptual proposition *p*, then *S*'s justification for believing *p* guarantees that *p*.

Although this chapter will be primarily expository, I will conclude by responding to the objection that two of the above theses—namely, McDowellian J-internalism and

²³ In employing the word 'dogmatism', I here allude to McDowell's tongue in cheek reply to Robert Brandom's essay, 'Knowledge and the Social Articulation of the Space of Reasons': 'What I urge in my paper is precisely that justification adequate to reveal a state as one of knowing must be incompatible with falsehood and can be had. ''Dogmatism''...is precisely what I defend.'(McDowell [2002], p. 98).

McDowellian C-externalism—are inconsistent with each other.²⁴ I argue that on McDowell's understanding of object-dependent singular thought, this objection fails.

3.2 McDowellian Justification Internalism

McDowell conceives of justification in terms of the Sellarsian metaphor of 'having the right standing in the space of reasons'.²⁵ This metaphor embodies a J-internalist intuition since it grounds knowledge in a subject's ability to present reasons for her beliefs, rather than in a causal (or some other non-inferential) connection between a belief and its object. In fact, the logical space of reasons is defined precisely in terms of that which stands in contrast to the logical space of lawful causal interactions. ²⁶ Central to McDowell's conception of the 'space of reasons' is the claim that all factors relevant to the epistemic standing of a subject's beliefs are subjectively available to her. Thus, McDowell asserts:

[O]ne's epistemic standing on some question cannot intelligibly be constituted, even in part, by matters blankly external to how it is with one subjectively. For how could such matters be other than beyond one's ken? And how could matters beyond one's ken make any difference to one's epistemic standing? (McDowell [1998], p. 390)

I interpret the locution 'how it is with one subjectively', as an umbrella term for the sorts of things that are typically taken to be internally available to one, such as one's thoughts, beliefs etc. By McDowell's lights, the circle delineating what is subjectively available to one exhausts that which may serve as a justifier for one's beliefs. When this idea is restated in the argot of possible worlds, we arrive at the following definition of McDowellian J-internalism:

²⁴ Significantly, to the extent that (M1) and (M2) are incompatible, so too are (M1) and (M4) since McDowell's claim that the justification of our perceptual beliefs is factive presupposes a content externalist framework.
²⁵ See: McDowell [1995].

²⁶ Thus, Wilfrid Sellars asserts: 'In characterizing an episode or a state as that of knowing, we are not giving an empirical description of that episode or state; we are placing it in the logical space of reasons, of justifying and being able to justify what one says' (Sellars [1963] p. 169).

McDowellian J-Internalism:

For all agents S_1 and S_2 and worlds W_1 and W_2 , if S_1 in W_1 and S_2 in W_2 are identical in terms of how things are with them subjectively, then S_1 and S_2 are identical in all respects relevant to the justification of their beliefs.

Under this conception, J-externalism can be described as the rejection of 'how things are with one subjectively' or 'the space of reasons' as the locus of epistemic justification:

According to a full-blown externalist approach, knowledge has nothing to do with positions in the space of reasons: knowledge is a state of the knower linked to the state of affairs known in such a way that the knower's being in that state is a reliable indicator that the state of affairs obtains. (McDowell [1995], p. 882).

McDowell shares the J-externalist intuition that there must be a non-accidental connection between a belief's justification and its truth. However, McDowell views the failure of traditional J-internalism to provide such a non-accidental connection as merely a symptom of a deeper problem—namely, 'the idea of the inner realm as self-standing, with everything within it arranged as it is independently of external circumstances' (McDowell [1986] p. 152). I will refer to McDowell's diagnosis of the source of traditional J-internalism's failure as the *interiorised conception*.

I take the interiorised conception to represent a marriage of J-internalism and Cinternalism. Jessica Brown [2004] defines C-internalism as the view that a subject's mental states are individuated wholly by her internal states, such as her brain states. On this view, the environment may cause a subject to be in a certain type of inner state and hence a certain mental state, but the environment is inessential to her being in that type of mental state. Once she is in the relevant type of brain state, she will also have the relevant mental state, regardless of what factors (i.e., the state of the world, the actions of a neuroscientist, or Cartesian demon) are responsible for her being in the particular type of brain state.²⁷ The upshot of C-internalism is that a subject's thought contents may be caused, but are in no way

²⁷ I owe the above description of C-internalism to Brown [2004], thought (following Tyler Burge) she opts for the label anti-individualism.

constituted by her environment. C-internalism is essentially a theory about mental content and how they are individuated; not an epistemological thesis. However, once it is combined with J-internalism, a gap is introduced between our beliefs and what makes them true. Admittedly, the advocate of the interiorized conception need not deny that our beliefs are causally influenced by the external world. However, the mere fact that a subject's thoughts are caused by objects in the external world, does not mean that her thoughts *represent* or are *about* those objects. Thus, on the interiorized conception, it remains un-established that a subject's thoughts are really *about* the world at all.

3.3 McDowellian Content Externalism

McDowell rejects the interiorised conception and favours a C-externalist account, according to which a subject's thought contents are not simply caused, but are partly constituted by objects in her environment. To wit, McDowell posits:

If we let there be quasi-Russellian singular propositions about, say, ordinary perceptible objects among the contents of inner space, we can no longer be regarding inner space as a locus of configurations which are self-standing, not beholden to external conditions; and there is now no question of a gulf, which it might be the task of philosophy to try to bridge, or declare unbridgeable, between the realm of subjectivity and the world of ordinary objects. We can make this vivid by saying, in a Russellian vein, that objects themselves can *figure in* thoughts which are among the contents of the mind. (McDowell [1986] p. 146)

The claim that 'objects themselves can figure in our thoughts' can be unpacked via an appeal to the notion of relational properties. Some property π is relational *iff* its possession by some object, α , requires the existence of some other object, β , to which α stands in a certain relation. For example, if α is smaller than β , then α 's property of being smaller than β is relational since its possession requires the existence of the contrasting (and in this case larger) object, β . By contrast, a non-relational or intrinsic property is one that can be possessed independently of the existence of other objects or events. An example of a non-relational property is water being chemically instantiated by H₂O.

On the interiorized conception, a subject's thought contents are wholly individuated by her intrinsic properties. Thus any two subjects that are identical in all their intrinsic properties also have the same thought contents. We may also conceive of the distinction between McDowell's C-externalism and the interiorized conception in terms of supervenience. Formally, X-group properties supervene on Y-group properties *iff* for all objects α and β , α and β cannot differ in their X-group properties without also differing in their Y-group properties. According to the interiorised conception, thought content supervenes solely on intrinsic properties²⁸. Thus, we arrive at:

The Interiorised Conception:

For all subjects S_1 and S_2 and worlds W_1 and W_2 , if S_1 in W_1 and S_2 in W_2 are identical in the intrinsic properties on which their thoughts supervene, then S_1 and S_2 are identical in all respects relevant to the justification of their beliefs.

By contrast, McDowellian C-externalism maintains that the contents of a subject's thoughts are partly individuated by objects in her environment.²⁹ Hence, two subjects, S_1 and S_2 , could be identical in all their intrinsic physical properties, and yet have different thought contents.³⁰ Thus, McDowellian C-externalism minimally entails the following claim:

McDowellian C-externalism:

There are subjects S_1 and S_2 and worlds W_1 and W_2 , such that S_1 in W_1 and S_2 in W_2 have the same intrinsic properties but differ in the content of their thoughts.

²⁸ There are numerous difficulties associated with the attempt to define an 'intrinsic property'. (See: Langton and Lewis [1998]). My use of 'intrinsic' here is purely a matter of convenience, but presumably these ideas can all be expressed without any appeal to the distinction between intrinsic and relational properties. However, if it were to turn out that the interiorised conception could not be articulated without appeal to the notion of intrinsic properties, then any difficulties confronting attempts to define intrinsic properties would only further impugn the interiorised conception.

²⁹ For examples of C-externalism that emphasize the object-dependence of singular thought see: Evans [1982]; Peacocke [1983]; and Campbell [1994].

³⁰ Alternative versions of C-externalism include Natural Kind C-externalism, according to which a subject's thought contents are individuated partly by the natural kinds in her environment (see: Putnam [1975b, 1975c] and Kripke [1980]) and Social C-externalism, according to which a subjects thought contents are partly individuated by the practice of her linguistic community (see: Burge [1979, 1986]).

The classic defence of C-externalism involves Twin-Earth arguments of the kind first suggested by Hilary Putnam [1975a, 1975b, 1975c] and later by Tyler Burge [1979, 1982, 1985, 1986, 1988]. Putnam invites us to imagine a subject, $S_{I_{i}}$ living in the actual world, $W_{I_{i}}$ who has had regular causal contact with water (i.e., H2O), but is ignorant of its chemical constitution. S_1 encounters a lake and has a thought that she expresses with the utterance, 'water is wet'. Let us call the mental property of being a thought with this content ' κ '. We may therefore say that S_1 's thought has κ . In the counterfactual world, W_2 , S_1 's psychophysical duplicate, S_{2r} also has a thought which she expresses with the utterance, 'water is wet'. However, on W_2 there is no H₂O but some qualitatively identical substance According to Putnam, the semantic content of the word 'water' that figures in S_i 's XYZ. utterance is different from that which figures in S_2 's utterance. Burge, building on Putnam's Twin-Earth example, argues that the subjects plausibly have different thoughts, since in S_1 's thought has κ_r , while S_2 's thought does not.³¹ Since variation in socio-environmental factors can make for variation in thought contents without variation in the intrinsic physical properties of the subject, it follows that a subject's thought contents fail to supervene on the her intrinsic physical properties. Thus, we arrive at C-externalism; the thesis that subjects may have the same intrinsic properties but differ in the content of their thoughts.

3.4 McDowellian Singular-thought Disjunctivism

McDowell is perhaps most famous (or is that infamous?) for his commitment to disjunctivism.³² The standard formulation of McDowell's disjunctivism is in terms of appearances:

[T]he essentially disjunctive conception of appearances...can allow what is given to experience in the two sorts of case to be the same in so far as it is an appearance that things are thus and so; that leaves it open that whereas in one kind of case what is given to experience is a mere appearance, in the other it is the fact itself made manifest. (McDowell [1998a], p. 389).

 ³¹ Cf. McGinn [1977].
 ³² For the traditional formulations of disjunctivism see J. M. Hinton [1973] and Paul Snowdon [1990].

However, an equally important, though typically less discussed aspect of McDowell's disjunctive conception is its application to singular thought:

There is a parallel contrast between two ways of conceiving singular thought: first, the idea that if one seems to be thinking about an ordinary external object in a way that depends on, say, its appearing to be perceptually present to one, the situation in one's inner world is either that one is entertaining an object-dependent proposition or that it merely appears that that is so. (McDowell [1986], p. 156).

We may gain a fuller appreciation of the construal of disjunctivism in terms of singular thought by contrasting McDowell's object-dependent picture of reference with Russell's Theory of Descriptions. According to the latter, thought connects with the object it is about via a description.³³ On this view, the singular component of a thought consists in a description, the φ , where the object referred to is the object that uniquely fits the description φ . McDowell summarises the upshot of such a picture:

The point of the alternative logical form proposed by the Theory of Descriptions is to ensure that the proposition that...a sentence is held to express is one available to be expressed in any case, whether or not there is something answering to the description. (McDowell [1986], p. 137).

According to McDowell's conception of object-dependent thought, one's thoughts connect with their objects not via a description but more directly in virtue of the relation in which one stands to those objects. On this view, the singular component of the thought is not equivalent to a description that remains constant whatever the state of the world. Rather, the content of the singular component of a thought is individuated partly by the object it is about. Thus, if one were counterfactually related to a different object, one would think a different thought. For example, suppose *S* is looking at a cube, α , and thinks the perceptual demonstrative thought 'that cube is red'. In the counterfactual situation in which *S* sees a different cube, β , she would have a different thought.

³³ See: Russell [1905].

Now consider a Twin Earth example where, instead of merely different objects, in one of the worlds there is no object at all. For example, suppose that while S_1 in W_1 is seeing and thinking about an actual red cube, S_2 in W_2 is merely hallucinating a red cube. On the descriptive account, S_1 and S_2 are having the same thought, the only difference between the two being the external fact that S_2 's thought lacks a reference. However, according to McDowell's object-dependent conception of singular thought, while S_1 is having a singular thought, S_2 fails to think any determinate thought at all. Like S_1 , she thinks she is seeing a red cube and having the demonstrative thought she takes herself to have. In short, she suffers an illusion of thought.³⁴ This of course does not mean that there is nothing going on in the subject's mind. Images, words and various subsidiary thoughts with genuine content may be taking place. But what is missing is a genuine singular thought with the content α is φ . Thus, we arrive at McDowell's singular-thought disjunctivism:

McDowellian S-disjunctivism

For any subject *S*, if it seems to *S* as if she is having the singular thought, α is φ , either the external object, α , figures in the content of her thought or she is suffering an illusion of thought.

The upshot of the disjunctive conception of singular thought is that veridical and hallucinatory visual experiences belong to different epistemic genera. Recall, that according to McDowellian J-internalism, only that which constitutes how things are with one subjectively, such as one's thoughts, beliefs etc, is relevant to one's epistemic standing. (This is a restatement of the familiar J-internalist slogan that only thoughts can act as reasons.)³⁵ In a case where one attempts to have an object-dependent singular thought, but fails to do so because the relevant object is missing, one *eo ipso* fails to have a reason. Thus, instances in

³⁴ See: Evans [1982], Boghossian [1997], and McDowell [1998b].

³⁵ I interpret McDowell's position here as analogous to the Davidsonian claim that 'only beliefs can justify beliefs' except that McDowell replaces 'beliefs' with 'thinkable content' (McDowell [1994], p. 28). Thus, McDowell insists that 'when we trace justifications back, the last thing we come to is still a thinkable content' (loc. cit.). However, McDowell rejects the coherentist conclusion that Davidson views as following from this claim. Instead, McDowell insists that 'experiences themselves are already equipped with conceptual content' (p. 25), which is to say, 'the world is made up of the sort of thing one can think' (Ibid, pp. 27-28). A detailed treatment of this idea falls outside the scope of the present study.

which one is entertaining a genuine singular thought have an epistemic value that is lacking in cases where one is suffering an illusion of thought. McDowell's disjunctive conception of object-dependent singular thought therefore translates into an epistemic disjunction between thoughts, which may constitute justifying reasons, and failed attempts at thought, which cannot.

3.5 McDowellian Justificatory Dogmatism

Finally, we come to McDowell's fourth thesis, which I have labelled *McDowellian J-dogmatism*. This position is instantiated in McDowell's assertion that statements of the form 'I see that...' provides one with 'an entitlement incompatible with any possibility of falsehood, to a claim whose content is given by the embedded proposition'. (McDowell [2002], p. 98). Essential to this thesis is the recognition that 'seeing' (in the epistemically relevant sense) is a success term. Along these lines, Alan Millar notes that '[t]here is a success element to seeing, or at least to seeing conceived in a very familiar way. If I see a cat, under this conception, then a cat must be there before me' (Millar [*forthcoming*]). To put the matter in McDowell's own words: 'Seeing that *p* constitutes falsehood-excluding justification for believing that *p*' (McDowell [2002], p. 97). The above ideas can all be summarised by saying that seeing is factive.

One way we may unpack the factivity of seeing is in terms of a *relational conception of perceptual experience*.³⁶ On this view, when a subject sees that p—for instance, that there is a vase on the table—she is having an experience that is essentially relational in that it would not be that very experience unless the relevant vase were there on the table. If *S* were to have a hallucination of a vase on the table, she would *eo ipso* be having a different type of experience. This does not mean that there could not be a hallucinatory experience of, say, a vase, which was indistinguishable from an authentic perceptual experience implicating a vase. However, experiences in which it appears to a subject that there is a vase before her may be described disjunctively:

³⁶An alternative way of unpacking the factivity of seeing is in terms of Millar's perceptual recognitional abilities. See: Millar [2007].

Experience Disjunctivism:

For any subject *S*, if it appears to *S* as if α is φ , either she is having an essentially relational experience implicating the external object α , or she is having a different type of experience.

The disjunctive conception of experiences allows us to preserve the desiderata that perceptual experiences are essentially relational, without denying the possibility of there being (phenomenologically) indistinguishable illusory experiences.

Let the locution 'seeing_{FACTIVE}' be a shorthand for the kind of seeing for which 'seeing that p entails p'.) 'Seeing_{FACTIVE}' is the instantiation of the propitious side of the disjunctive conception of perceptual experiences, and represents one upshot of such a conception. Seeing_{FACTIVE} that p provides one with justification for believing that p since seeing_{FACTIVE} that p puts one in the position to have the relevant object-dependent singular thought. Moreover, as we noted above, having a singular thought means that the object, α , that the thought is about, figures in the content of one's thoughts, which just is to fulfil the justificatory requirements of McDowellian J-internalism. Since seeing_{FACTIVE} that p excludes the possibility that $\sim p$, and given that seeing_{FACTIVE} that p justifies one's belief that p, that which justifies one's belief that p excludes the possibility that $\sim p$. Thus, we arrive at McDowellian Dogmatism:

McDowellian J-dogmatism:

For any subject *S*, if *S* has perceptual knowledge that *p*, then *S*'s justification for believing *p* guarantees that *p*.

We can now see how McDowell's four-fold approach allows us to respond to the perceptual Gettier cases presented in §2. Recall the Gettier case of veridical hallucination described earlier: *S* hallucinates that there is a vase on the table, which serves as the basis for her belief that (a), 'there is a vase on the table'. As luck would have it, there *is* a vase on the table (though it is obscured by some books and has no part in *S* coming to believe that there
is a red cube on the table). By McDowell's lights, all of the following are true in the above case:

- (M5) In both GTTR_{VH} and GTTR_{LR}, given that the (actual) vase falls outside S's visual field, *S* does not see the vase in the epistemically relevant sense of seeing_{FACTIVE}, *a la* McDowellian J-dogmatism.
- (M6) In both GTTR_{VH} and GTTR_{LR}, given that S does not see the vase in the epistemically relevant sense of seeing_{FACTIVE}, the vase does not figure in *S*'s thought contents, *a la* McDowellian C-externalism.
- (M7) In both GTTR_{VH} and GTTR_{LR}, given that the vase does not figure in *S*'s thought contents, *S* fails to have the relevant singular thought, *a la* McDowellian S-disjunctivism.
- (M8) In both GTTR_{VH} and GTTR_{LR}, given that *S* fails to have the relevant singular thought, *S* does not have a reason to believe that (a), and her belief is therefore unjustified, *a la* McDowellian J-internalism.

The upshot of (M5)-(M8) is that in both GTTR_{VH} and GTTR_{LR}, one may be described as having the true (existentially quantified) belief that *there is a vase on the table*.³⁷ But one does not have a *justified* belief that *there is a vase on the table*. This strategy applies with equal force to the Gettier case against local reliabilism.

³⁷ Note that though one lacks the relevant singular thought, one may still have the existentially quantified thought. However, with the singular component missing, the existentially quantified thought remains without basis and therefore unjustified. The nature of the relationship between existentially quantified and singular thoughts remains a matter of significant debate. For example, Harold Noonan [1991] argues that singular thoughts are explanatorily redundant and even suggests that they may not exist. For a reply to Noonan, and a defence of the claim that singular thoughts play an essential explanatory role, see Sawyer [2003] and [2004]. I will not attempt to settle this debate here.

3.6 McDowellian J-Internalism and C-Externalism Reconciled³⁸

A distinctive feature of McDowell's response to the perceptual Gettier limned above is that it combines the C-externalist claim that one's thoughts are partly individuated by objects in one's environment with the J-internalist claim that all factors relevant to the justification of one's beliefs are subjectively available to one. But some have insisted that J-internalism is inconsistent with C-externalism.³⁹ The reason for this can be seen easily enough when we conceive of J-internalism in terms of the interiorised conception. Recall, the interiorised conception (IC) and C-externalism (CE) amount to the following claims:

- (IC) For all subjects S_1 and S_2 and worlds W_1 and W_2 , if S_1 in W_1 and S_2 in W_2 are identical in the intrinsic properties on which their thoughts supervene, then S_1 and S_2 are identical in all respects relevant to the justification of their beliefs.
- (CE) There are subjects S_1 and S_2 and worlds W_1 and W_2 , such that S_1 in W_1 and S_2 in W_2 have the same intrinsic properties but differ in the content of their thoughts.

There is a conflict between (IC) and (CE) since the former requires all subjects who are identical in terms of their intrinsic properties to have the same justificatory properties, while the latter allows subjects with identical intrinsic properties to differ with regards to the justificatory properties of their beliefs. For example, consider two subjects, S_1 and S_2 , who are identical in terms of their intrinsic properties, but occupy different external environments, W_1 and W_2 , respectively. Suppose S_1 and S_2 both performed the following valid deduction:

³⁸ The present defence of the compatibility of McDowellian J-internalism and C-externalism builds on the compatibilist arguments of Brueckner [2002].

³⁹ The claim that C-externalism is inconsistent with J-internalism, or the so-called 'incompatibility thesis', was first alluded to by Paul Boghossian. However, Boghossian only discusses the incompatibility thesis indirectly, in the context of discussions of authoritative self-knowledge (see: Boghossian [1989] and [1997]). For an explicit defence of the 'incompatibility thesis' see BonJour [1992] pp. 132-136, and Vahid [2003]. For criticisms of the 'incompatibility thesis' see Chase [2001] and Brueckner [2002]. For arguments in support of the claim that McDowell's C-externalism, in particular, entails J-externalism, see Greco [2004].

- (A) Water is a liquid.
- (B) Water is potable.
- (C) Therefore, water is a potable liquid.

According to (CE), the thoughts expressed by sentences (A)–(C) are different for S_1 and S_2 . This is because S_1 's thoughts are individuated in terms of water while S_2 's thoughts are individuated in terms of twater. The distinct pairs of thought that S_1 and S_2 express by (A) and (B) are relevant to the justification of the pair of beliefs they respectively express by (C). Consequently, S_1 and S_2 satisfy the antecedent of (IC), since *ex hypothesi* they are identical in terms of their intrinsic properties, but fail to satisfy the consequent—to wit, they differ in some respects relevant to the justification of their beliefs. Thus, if (CE) is true, then (IC) must be false.

There is a strong temptation to assume that this line of argument also impugns McDowellian J-internalism, but this temptation should be resisted. Recall, McDowellian Jinternalism (MI) amounts to the following claim:

(MI) For all subjects S_1 and S_2 and worlds W_1 and W_2 , if S_1 in W_1 and S_2 in W_2 are identical in terms of how things are with them subjectively, then S_1 and S_2 are identical in all respects relevant to the justification of their beliefs.

If we take seriously McDowell's notion of object-dependent thought there is no obvious inconsistency between (MI) and (CE). According to McDowell, how it is with one subjectively is in part constituted by objects in one's environment. For instance, in the Twin Earth examples generated by (CE), the object-dependent thoughts expressed by (A) and (B) in the foregoing deductive inference are different for S_1 and S_2 . Hence, by McDowell's lights, S_1 and S_2 fail to satisfy the antecedent of (MI) since S_1 and S_2 are *not* identical with regards to how things are with them subjectively. Thus, McDowell can, without contradiction, continue to hold to (MI) while subscribing to (CE).

3.7 Chapter Summary

In this chapter I have argued that McDowell's account of perceptual knowledge should be interpreted as an attempt to preserve the J-internalist intuition that one's beliefs can only be justified by factors that constitute how things are with one subjectively and the J-externalist intuition that there must be a non-accidental connection between justification and truth. I have used the Gettier case of veridical hallucination as a foil for bringing out four theses McDowell employs towards this end. Moreover, I have argued that the assumption that McDowell's commitment to C-externalism is inconsistent with his version of J-internalism is unwarranted. Consequently, I believe McDowell's conception of perceptual knowledge represents the first step towards a post-Gettier conceptual analysis of knowledge. In the next chapter I will examine the objection that McDowellian C-externalism is incompatible with authoritative self-knowledge.

4. SINGULAR THOUGHT EXTERNALISM AND SELF-KNOWLEDGE

At least since Descartes, most philosophers have believed that self-knowledge is different, in significant ways, from knowledge of the external world and knowledge of other people's thoughts. Specifically, it is widely assumed that we have special access to our own thoughts in contradistinction to the thoughts of others and facts about the external world. Thus, there is taken to be a fundamental asymmetry between self-knowledge and other forms of knowledge. We may sum up the above intuitions by saying that self-knowledge is *authoritative*.

In her recent paper, 'Self-Knowledge and Inner Space', Cynthia MacDonald argues that McDowell's brand of singular thought C-externalism is incompatible with authoritative self-knowledge.⁴⁰ Recall, according to singular thought C-externalism, if one attempts to think a thought of the form α is φ in a situation in which the relevant object, α , is missing, one fails to have the relevant thought. Since one may be unaware of this fact, it is possible for one to think that one is entertaining a certain thought at a certain region of one's mind when there is in fact only a gap. In such a case, one suffers an illusion of thought. MacDonald thinks that this fact renders singular-thought C-externalism incompatible with authoritative self-knowledge. She summarises the problem as follows:

What problem does [singular-thought] externalism seem to pose for the view that subjects have authoritative self-knowledge? It seems to imply that subjects can be mistaken not only about whether the thoughts they are actually thinking are true of the world around them, but also abut whether they are thinking certain thoughts it seems to them that they are thinking. More generally, it seems to imply that others can be better placed than a subject to know whether that subject is thinking thoughts of the *kind* she takes herself to be thinking, irrespective of whether she *is* thinking those thoughts. (MacDonald [2006] p. 74)

⁴⁰ The expression, 'authoritative self-knowledge' is a term of art and has been used in different ways by different philosophers. However, since both MacDonald and McDowell agree with the asymmetry thesis, my definition of authoritative self-knowledge represents common ground between the two. (See: MacDonald [2006], p. 73ff; and McDowell [2006], p. 92-94).

It is important to get clear on what MacDonald is asserting in this passage. Two points are worthy of mention. First, as can be seen from her concluding sentence, MacDonald is not only suggesting that singular-thought C-externalism entails that another may be better placed regarding whether or not a subject is thinking a particular thought. What MacDonald takes to be most problematic about singular-thought C-externalism is that it allows another to be better placed to know what *kind* of thought a subject is having. This is significant since, as we shall see below, McDowell concedes the former point—that another may be better placed to know what kind or type of thought one is having. It is the latter claim, then, which is the point of contention between the two thinkers. I will also attempt to show that it is on the latter claim that debate over the compatibility of singular-thought C-externalism and authoritative self-knowledge hangs.

Secondly, it should be observed that MacDonald advances two separate criticisms against singular-thought C-externalism:

- Singular-thought C-externalism implies that a subject can be mistaken about her thoughts (in ways to be spelt out below).
- (2) Singular-thought C-externalism implies that another may be better placed than a subject with regards to her thoughts (in ways to be spelt out below).

In so doing, MacDonald conflates two separate claims about authoritative self-knowledge that I wish to keep separate—namely, that authoritative self-knowledge is (1) infallible and (2) superior to knowledge of other minds and the external world. MacDonald does not distinguish between these two understandings of self-knowledge, which I believe sometimes makes her argument less clear than it could be. In order to remedy this problem I will deal with the issue of infallibility and superiority separately. In this chapter I will be arguing that MacDonald has given us no reason to think that singular-thought C-externalism is incompatible with authoritative self-knowledge.

4.1 Infallibility and Authoritative Self-Knowledge

In its most extreme form the label 'infallible self-knowledge' may be applied to the claim that a subject cannot be mistaken about the *type* or *content* of her thoughts. This claim is often referred to by McDowell as the 'Cartesian view'. Although often attributed to Descartes, it is almost certainly not a position he ever endorsed. First, textual evidence suggests that Descartes was actually primarily concerned with indubitability rather than infallibility.⁴¹ Recall, Descartes defines knowledge in terms of doubt. The requirement that knowledge be based on complete certainty amounts to requiring a complete absence of doubt or an 'inability to undermine one's conviction'. Given the Cartesian construal of knowledge, it may therefore be argued that infallibility never, strictly speaking, enters into the picture.⁴² Second, Descartes only reserved the label of 'complete certainty' for a subset of our selfknowledge – namely, those that are 'clear and distinct.' Thus, there is no reason to attribute to him the idea that infallibility is a characteristic of self-knowledge *simpliciter*. Even so, if pressed on the matter, it seems likely that Descartes would hold that at least some of our selfknowledge is infallible and, at the very least, there is one thought about which we could never be mistaken—namely, the *cogito*. Simply put, the *cogito* refers to Descartes' famous, 'I think', or better 'I am now thinking'. This thought is self-verifying since the very act of thinking it makes it true. Tyler Burge [1988] adds to the original *cogito*-thought, thoughts of the form, 'I think (with this very thought) that writing requires concentration' and 'I judge (or doubt) that water is more common than mercury'.43

The notion of cogito-thoughts presents us with a moderate version of infallible selfknowledge. On this view, we may be mistaken about *some* of our thought contents. However there is at least one domain—namely, thoughts which take themselves as their objects—about which we are infallible. Burgean C-externalism already comes well-equipped to preserve the intuition that we are infallible with regards to such thoughts. According to

⁴² I mention this here, not in an attempt to settle a historical question but rather to begin loosening an idea that is often taken for granted—namely, that infallibility is fundamental to the concept of self-knowledge.

⁴¹ See: *Meditations on First Philosophy*.

⁴³ Notice that the objects of these self-ascribed thoughts are standing propositions and it doesn't matter if the proposition is wrong or right. The question is whether I am thinking a thought with a certain content, not whether the content thought about is true. We shall soon see that this does not extend to cases of object-dependent singular thought.

Burge, cogito-thoughts have a reflexive, self-referential character. When one thinks that 'I am thinking that p', one not only has the first-order thought that p but also thinks about it as one's own. Since the object of the second-order judgement is none other than the first-order content, it is impossible to misjudge the content of the first order thought (even if this content partly depends on external conditions). For clarity, I will refer to Descartes' original cogito-thought as the Cartesian cogito-thought and Burge's additional cases as Burgean cogito-thoughts. According to the moderate understanding of infallible self-knowledge, a subject may be mistaken about some of her mental contents but not with regards to her Cartesian or Burgean cogito-thoughts.

4.1.1 The Illusion of Thought Objection

Significantly, the infallibility that characterises Burgean cogito-thoughts does not extend to object dependent thoughts à la singular-thought C-externalism. To see this we must distinguish between two types of cases, slow switch cases in which a subject is transported from a world in which there is only water to one in which there is only twater, and cases in which the subject takes herself to be looking at water when there is only a mirage, or some other visual illusion, before her. In the former case, singular-thought Cexternalism is well placed to exploit the notion of self-verifying cogito-thoughts. This is because singular-thought C-externalism agrees with Burgean C-externalism regarding how such cases should be treated. However, in the latter case singular-thought C-externalism states that the subject in question does not have a thought at all. The subject takes herself to be entertaining a thought of the relevant kind when there is no such thought available to be had. Recall, according to Burgean C-externalism, the second-order thought takes the content of the first-order thought as its object. But according to singular-thought C-externalism, the first-order thought lacks any content, and there is therefore nothing for the second-order thought to take as its object. It would therefore seem that to the extent to which the subject is mistaken with regards to the first order thought, she is also mistaken with regards to the second order thought which takes the content of the first order thought as its object.

4.1.2 The 'Truth-Gap' Argument

There is a tempting reply to the above criticism that, though ultimately unsuccessful, serves to further clarify and underscore the extent to which singular-thought C-externalism undermines the infallibility of Burgean cogito-thoughts. The reply I will now consider is rooted in the idea that when a thought is not available, because it lacks a denotation, the result is not a false belief but rather a truth-value gap. A thought of the form α is φ is false when its negation, α is not φ , is true. However, if the case of a missing denotation were to be considered false, then this would mean that it also entailed the negation of α is φ . However, this would be to fail to distinguish between a proposition's truth-value and what the proposition presupposes (i.e., what circumstances must obtain for a sentence to have a determinant truth-value). McDowell arrives at a similar conclusion:

The syntax of sentences of the relevant sort fits them to express singular thoughts if any; where a singular thought is a thought that would not be available to be thought or expressed if the relevant object, or objects, did not exist. It follows that if one utters a sentence of the relevant sort, containing a singular term that, in that utterance, lacks a denotation, then one expresses no thought at all; consequently, neither a truth nor a falsehood. (McDowell [truth value gaps] p. 204).

Analogously, it may be argued that the case in which a subject suffers an illusion of thought does not undermine her infallibility since, strictly speaking, she does not have a false cogito-thought. In fact, she fails to have a cogito-thought at all. Instead, there is only a gap in the inner configuration of the subject's mind to which neither the label of true nor false can be applied. Admittedly, there are other thoughts in the vicinity about which the subject is in fact mistaken. However, this would only count against singular-thought C-externalism if we buy into the implausible view that infallibility is a characteristic of self-knowledge *simpliciter*. But short of the extreme view, the infallibility of cogito-thoughts is still preserved vis-à-vis singular-thought C-externalism.

Unfortunately, the 'Truth-value Gap Argument' misses its mark. According to Burge, cogito-thoughts take the form 'I am thinking (with this very thought) that p', where p is some proposition of the form ' α is φ' . According to singular-thought C-externalism, if the

object in question, α , is missing, then the singular expression ' α is φ ' is no longer available to be thought. Thus, for the cogito thought 'I am thinking (with this very thought) that p' there is only a (truth-value) gap where p should be. But notice, the truth-value gap only applies to p itself—to wit, the expression ' α is φ ' is neither true nor false. However, the sentence in which p figures, 'I am thinking (with this very thought) that p' still has a determinate truthvalue. Since p is empty or missing, the sentence 'I am thinking (with this very thought) that p' is plainly false. Thus, given singular-thought C-externalism, a subject may even be mistaken about her Burgean cogito-thoughts.⁴⁴

Let us take stock of what we have established thus far. Both Burgean and singularthought C-externalism are consistent with the infallibility of the Cartesian cogito thought. Consequently, even on the McDowellian 'anti-Cartesian' conception of the mental, a subject may be infallible regarding the fact that she is now thinking. However, singular-thoughts that take the form of Burgean cogito-thoughts are not infallible. Given that singular-thought C-externalism does not preserve the infallibility of Burgean cogito thoughts it is not compatible with moderate infallible self-knowledge.

Does the above conclusion pose a problem for the compatibility of singular-thought C-externalism and authoritative self-knowledge? I do not believe it does. To see why, we may distinguish between self-knowledge *simpliciter* and what Burge refers to as 'basic self-knowledge'.⁴⁵ The latter refers to cogito-thoughts which are reflexive by nature, and therefore self-verifying. The former include our self-ascriptions of standing propositional attitudes. Examples include judgements such as 'I believe that invading Iraq was a mistake', or 'I regret voting for Bush,' where each respective judgement is about a standing state, the self-attribution of which is triggered by a particular occurrence of the thought. Such instances of self-knowledge are arguably much more prevalent than so-called cogito-thoughts. Burge acknowledges that such self-ascriptions may still constitute self-knowledge even though they are not self-referential or self-verifying.⁴⁶ Thus, Burge maintains that the

⁴⁴ This does not, however, have any bearing on Descartes' original cogito-thoughts, 'I think' and 'I am now thinking' since Cartesian cogito-thoughts are not object dependent in the way that singular-thoughts are. ⁴⁵ See: Burge [1988], p. 658.

⁴⁶ Along these lines, Burge notes that reflection on cogito thoughts yield, 'at most an illuminating paradigm for understanding a significant range of phenomena that count as self-knowledge....A full discussion of authoritative self-knowledge must explicate our special authority, or epistemic right, even in numerous cases where our judgements are not self-verifying or immune to error' (Burge [1988], p. 658).

self-verifying character of cogito thoughts cannot be generalised to all forms of selfknowledge. However, although he acknowledges that not all forms of self-knowledge are self-verifying, he still insists that all cases of authoritative knowledge are (in the case of a rationally competent subject) infallible.

I think that, in all cases of authoritative knowledge, brute mistakes are impossible. All errors in matters where people have special authority about themselves are errors which indicate something wrong with the thinker. Dealing with the whole range requires subtlety. But the point as applied to what I take to be the basic cases is straightforward. No errors at all are possible in strict cogito judgments; they are self-verifying. (Burge [1988], p. 658).

It is at this point that McDowell and Burge part ways. McDowell agrees with Burge that there are certain basic forms of self-knowledge about which a subject is infallible. For McDowell, these include the original Cartesian cogito, Burgean cogito-thoughts which do not include singular propositions and our knowledge of how things seem to us. However, while Burge maintains that infallibility also characterises those forms of self-knowledge which are not self-verifying, McDowell limits infallibility to non-singular cogito-thoughts and how things seem to a subject.

The key to understanding McDowell's position on the question of infallibility and self-knowledge lays in his parallel treatment of the disjunctive conception of appearances and the disjunctive conception of singular thought. According to the disjunctive conception of appearances, perceptual experiences are either cases of the fact that p making itself perceptual manifest, or it only appearing as if that is the case. A subject may be mistaken about which of the two disjuncts obtains, but not with regards to the fact that she is having an appearance as if p. This matches the putative intuition, or what McDowell refers to as the 'ancient idea', that 'appearances are not open to question.'⁴⁷ Thus, McDowell maintains that a subject is infallible about how the external world appears to her. He then proceeds to 'exploit an analogy between the disjunctive conception of perceptual appearances and a disjunctive conception of situations in which it at least seems to one that one is entertaining an object-dependent thought.' As with the case of perceptual appearances one is infallible with

⁴⁷ See: McDowell [2006], p. 91. See also McDowell [singular].

regards to what one seems to be thinking. Moreover, like the perceptual case, seeming to have a singular thought is construed by McDowell as having disjunctive truth-conditions. It may be made true by one actually having a singular thought because the object the thought is about figures in the thought, or as a case in which one merely seems to do so but does not because the relevant object is missing. Significantly, thoughts of the form, 'I seem to be thinking that p' are not themselves object dependent thoughts, though they may take object-dependent thoughts as their object.⁴⁸ That is, if one had the thought, 'I seem to be thinking that that tiger has stripes', the second order thought is not only true in cases where one is actually have a contentful thought about a tiger. In cases in which one fails to have a contentful tiger thought. Since the second-order thought is true whether or not there is a tiger content present, one cannot be mistaken about how things seem to one.

McDowell refers to the idea that appearances are not open to question, and the corresponding idea that one may be infallible regarding what one seems to be thinking as the less than fully Cartesian view (henceforth the *quasi-Cartesian view*). The quasi-Cartesian view agrees with the 'ancient idea' that a subject cannot be mistaken about how things appear to her. However, the ancient view did not see 'appearances' themselves as facts with a corresponding truth value and so fails to take the extra step we find in the quasi-Cartesian view. Namely, the quasi-Cartesian view extends the concept of truth to how things appear to a subject, and therefore makes the way things seem a fact about which one may never be mistaken.

In order to arrive at the fully Cartesian view, McDowell points out that we must take an additional step.

We need to suppose that the 'inner' region of reality that is newly recognised with the first step is exhausted by the facts brought into view by that step: facts infallibly knowable by the subject, to the effect that things appear a certain way to her. Only with this second step do we arrive at the picture of a region of reality whose layout is through and through available to an infallible cognitive capacity. (McDowell [2006], p. 90).

 $^{^{48}}$ This point is akin to that made earlier in response to the truth-gap argument. Although the content of the singular thought, α is φ , lacks a truth value due to a missing denotation, the proposition that the singular statement figures in still has a determinate truth-value.

It is the fully Cartesian view (though, as I pointed out above, Descartes himself may not have held it) that drives the idea that authoritative self-knowledge must be infallible. Once we reject the fully Cartesian view, we are free to embrace the idea that there are segments of a subject's mental life about which she may be mistaken. Thus, while a subject may be infallible with regards to the fact that she *seems* to be thinking a thought with a certain content, she is not infallible regarding the fact that she *is* thinking that particular thought content. A subject may be infallible with regards to the fact, construed disjunctively, that she seems to be thinking a certain thought; to wit, that she is either having a singular thought or she appears to do so. But she is not infallible about which of the two disjuncts obtains.

4.2 The First MacDonald Objection

According to MacDonald, a subject is *authoritative* with regards to some thought θ iff she is 'better placed' to know that she is having θ than anyone else. Thus understood, she argues that authoritative self-knowledge is incompatible with singular-thought C-externalism. MacDonald advances three arguments on this score. Her first objection runs as follows:

The first concerns situations in which a subject may take herself to be entertaining a thought of a certain kind but is not, because to do so would require the presence of a particular object in her visual field where none exists. Consider, for example, a subject who lives in an environment in which there are tigers and who attempts to think a thought of the type 'that tiger has stripes' in a situation in which there is no suitably situated tiger. Although she may have thought token thoughts of that type in the past in contexts in which tigers were visually present to be demonstrated, she fails in this case to think a token thought of that type. This kind of cases threatens authoritative self-knowledge because it is one where the subject mistakenly takes herself to be thinking a thought that is not, in her situation, available for her to think; and another may be better placed than her to know this. (MacDonald [2006], pp. 74-75).

Earlier, we noted that McDowell subscribes to the infallibility of only some types of self-knowledge, such as the knowledge of what thought content one seems to be thinking. By his lights, other forms of self-knowledge, though authoritative, are not infallible. Similarly, McDowell may argue that a subject is always 'better placed' to know what type of content she seems to be thinking, but not whether she is actually thinking such content. In fact, McDowell appears to make precisely this concession:

[T]here is nothing ontologically or epistemologically dramatic about the authority which it is natural to accord to a person about how things seem to him. This authority is consistent with the interpenetration of the inner and outer, which makes it possible for you to know the layout of my subjectivity better than I do in a certain respect, if you know which of those two disjuncts obtains and I do not. In this framework, the authority which my capacity for 'introspective knowledge secures for me cannot seem to threaten the very possibility of access on your part to the facts within its scope. (McDowell [1986] p. 154)

However, the fact that sometimes another may be 'better placed' than one to know whether or not one is having a certain thought content falls short of another knowing how things seem to one. At this point, McDowell may invoke the essentially disjunctive character of singular thought. When it seems to *S* as if she is thinking a thought of the form α is φ , this seeming is a case of either: the subject actually thinking a thought of the form α is φ , or her suffering the illusion of entertaining a thought of this kind. Another may be better placed to know which of the two disjuncts obtains since she may know whether the object which settles the disjunction is present or not. However, the subject is still better placed to know how things seem to her—i.e., that the disjunction obtains at all. Thus, there is no conflict between what another is better placed to know and what the subject herself is better placed to know *which* disjunction obtains while the subject is better placed to know *whether* the disjunction obtains.⁴⁹ Thus, McDowell already seems equipped with a reply to MacDonald's first objection.

⁴⁹ Some would no doubt consider this already too great a concession on McDowell's part since it already appears to place singular-thought C-externalism at odds with our putative understanding of authoritative self-knowledge.

The above reply to MacDonald's first objection serves to bring into sharp focus the point on which the thinkers disagree. McDowell agrees with MacDonald that another may be better placed with regard to at least one aspect of our self-knowledge—namely, which of the two disjuncts obtains. However, he insists that we may still preserve the intuition that we are always better placed with regards to how things seem to us. MacDonald disagrees, insisting that on McDowell's account a subject may even be mistaken about how things seem to her. MacDonald's argument on this score is most clearly seen in her second objection.

4.3 The Second MacDonald Objection

MacDonald's second objection attempts to fill in the loop-hole left by the first objection. Here is the objection in her own words:

Suppose that a subject is attempting to think a thought of the type 'that tiger has stripes', not just in a situation in which there is no suitably situated tiger present in her visual field, but in a world (Twin Earth) in which there are no tigers at all but only pligers—creatures which look, feel, and behave like tigers but have a different biological constitution. In this case there is no tiger content—real or apparent—in the inner realm to constitute the content of her thought. In this case, unlike the first one, it is not even possible for the subject to entertain singular token tiger thoughts on other occasions; and another, who has better knowledge of the facts of biology and knows that the world in which they are living is not Earth but Twin Earth, might be in a better position than the subject is to know this. (MacDonald [2006], p. 75).

One prima facie objection to MacDonald's argument runs as follows. MacDonald takes it to be significant that while the subject described in her first objection is unable to have a certain token thought, the subject described in the second objection is unable to have any thoughts of a certain type. However, it is not immediately clear why this should make any difference vis-à-vis singular-thought C-externalism. Recall that according to singular-thought C-externalism, if a subject takes herself to be thinking a thought of the form α is φ , in a situation in which there is no object, α , present, the subject suffers an illusion of thought.

However, in §4.5 I argue that even with this concession, singular-thought C-externalism still preserves what is essential to authoritative self-knowledge.

The same is equally true of cases where a subject takes herself to be thinking a thought of the form α is φ in a world without any α s. Just as in the first case, the subject fails to think a thought of the type she takes herself to think and therefore suffers an illusion of thought. The fact that in the first case she happens to be living in a world with tigers while in the second she happens to be in a world without tigers makes no difference. Since McDowellian C-externalism, which is concerned with singular thought, treats each perceptual situation on a case by case basis, the distinction between token and type thoughts is not salient. Thus, the second MacDonald objection seems to simply collapse into the first. If the above argument is right, then the reply to the first objection offered in §4.2, is equally effective against MacDonald's second objection.

However, I believe there is an alternative and much more charitable way of reading MacDonald's second objection. Recall, the reply to MacDonald's first objection presented in §4.2 attempts to preserve the intuition that a subject is better placed with regards to at least some of her thoughts by arguing that what a subject may be authoritative about in this regard is different from what another my be authoritative about. Specifically, the subject is better placed to know how things seem to her—that she seems to be having a thought of the form α is φ ; that is, she is better placed to know that her thought is of a certain type. MacDonald's second objection attempts to impugn the above claim. In the case of a subject that has been unwittingly transported to Twin Earth, there are no α -type thoughts available to be thought at all. However, since the subject is unaware of her switch between worlds, someone else who is aware that she is now on Twin Earth may be better placed to know what type of thought she seems to be having. Thus, MacDonald's second objection seems to represent a case where another is better placed, not only with regards to which of the two disjuncts holds, but with regards to how things seem to the subject.

Unfortunately, MacDonald's second objection—when charitably read—still rests on a questionable assumption. Namely, the claim that the subject in question cannot have α -type thoughts. Recall, MacDonald invites us to imagine a subject who attempts to think tiger-type thoughts in a world only inhabited by pligers. She maintains that there are simply no tiger-type thoughts available for the switched subject to think. The only way to make sense of the switched subject attempting to think tiger thoughts would be to imagine a subject who lived

on Earth, thereby acquiring the concept of tiger, and who is then transported to Twin-Earth. But, as McDowell points out, we have no reason to assume that such a subject would be unable to think tiger-type thoughts.⁵⁰ Of course, the subject's tiger-type thoughts would fail to apply to any of the creatures in her present environment. Thus, if she were to point to one of these creatures and think to herself, 'oh, that's a tiger', she would be mistaken. However, it does not follow from this that she would be unable to think tiger-type thoughts or that she would be unaware of what type of thought she *seemed* to be thinking. The fact remains that the switched subject would still have the tiger concept she acquired when she was on Earth.

There are two separate but closely related questions that bear on the above reply to MacDonald's second objection.

- (Q1) In what sense can one have a tiger-type thought if one is in a *situation* in which there is no tiger content available?
- (Q2) In what sense can one have a tiger-type thought if one is in a *world* in which there is no tiger content available?

McDowell's answer to the first question is suggested by the following passage:

Particular *de re* senses, each specific to its *res*, can be grouped into sorts. Different *de re* senses (modes of presentation) can present their different *res* in the same sort of way: for instance, by exploiting their perceptual presence. And the univocity of a context-sensitive expression can be registered by associating it with a single sort of *de re* sense Given a context, a sort of *de re* sense may determine a *de re* sense (if one cares to put it like that), or else it . . . may determine nothing. And in the latter sort of case, according to this way of thinking, there can only be a gap—an absence—at, so to speak, the relevant place in the mind—the place where, given that the sort of *de re* sense in question appears to be instantiated, there appears to be a specific *de re* sense. (McDowell [De Re sense], pp. 220-221).

⁵⁰ See: McDowell [2006], p. 93.

Thus, McDowell maintains that in the case in which a subject's thought lacks content, we can determine what type of thought the subject seems to have based on the sort of *de re* sense that appears to be instantiated in the particular place in the subject's mind. MacDonald seems to concede this very point:

One might say that what gives the appearance to a subject that she is thinking a particular token tiger thought is precisely that there is available to her a sort of de re sense, the tiger sort, which she mistakenly thinks in this situation is instantiated in a particular place in her mind. (MacDonald [2006], p. 80, Italics hers).

One way we may cash out the above idea is by distinguishing between a thought and a thought-sign (a thought being a thought-sign that has content).⁵¹ On this model, the difference between the case in which one has an object dependent thought and one in which one does not is that in the former, the thought-sign in question has a content while in the latter, it does not. If we accept this suggestion, a thought-sign should not itself be seen as 'an aspect or ingredient of content' but rather as 'a putative bearer or vehicle of content'.52 Alternatively we may unpack this idea in terms of McDowell's calibrated understanding of Fregean 'mock thoughts'. According to this proposal, the no-reference situation in which a subject takes herself to be thinking about a tiger she suffers a mock tiger-thought, the crucial point being that although there is no genuine tiger thought, the mock thought is one that has to do with tigers rather than pligers. In response to (Q1) we may say, therefore, that in a noreference situation in which there is no tiger content present, the subject has a mock tigerthought.

As we noted above, MacDonald agrees with McDowell's reply to (Q1) since she grants that a subject may have a tiger-sort of *de re* sense in a *situation* in which there is no tiger. Where she and McDowell part ways however is on whether a subject can have a tigertype *de re* sense in a *world* in which there are no tigers. On this question, MacDonald's emphatic answer is no:

⁵¹ Cf. Carruthers [1987], p. 20.
⁵² See: McDowell [De Re], footnote 19.

Suppose that I am attempting to think a thought of the type, that tiger has stripes, not just in a situation in which there is no suitably situated tiger present in my visual field, but in a world in which there are no tigers at all but only pligers. In this case there is no tiger content—real or apparent—in the inner realm to constitute either disjunct (1) or disjunct (2). Not only is there no particular *de re* sense, but there is no sort of *de re* sense. (MacDonald [2006], p. 81).

And here is precisely where MacDonald errs. She mistakenly assumes that since there are no tigers on Twin Earth, singular-thought C-externalism implies that there are no tiger-type thoughts available. Of course, ordinarily this would be the case since a subject who only lived on Twin Earth, and never stood in the appropriate relation to tigers, would be unable to have tiger-type *de re* senses. In the normal Twin Earth case, MacDonald points out that the problem is not simply that the Twin-Earther fails to have a token tiger thought on the occasion in question, but that she is unable to have it on any other occasion. But this is not a problem with the switched subject who, though now on Twin Earth, was able to have tiger de re senses of the tiger sort, albeit unsuccessfully for the duration of her time on Twin Earth.⁵³ In fact, this is the only way we could make sense of the intuition that the subject makes a mistake when she thinks that the striped animal before her is the same as the striped animal she saw and/or learned about on Earth.

4.4 The Third MacDonald Objection

MacDonald's third objection describes a case where another seems better placed than a subject, not because the subject is in some way mistaken about her thoughts, but because another subject seems better placed to tell that her thoughts really are of the kind she takes them to be. Again, here is the objection in MacDonald's own words:

⁵³ Strictly speaking, this last point is not quite right since, like more standard forms of C-externalism, singularthought C-externalism allows that a subject who remained on Twin-Earth for a sufficient length of time would eventually begin to think pliger thoughts. However, this does not make a difference to the present line of argument since in a slow switch case the subject's thought would match the animal in question so that the scenario MacDonald envisions never arises.

Thus, consider a situation in which someone knows the facts about biology and Twin Earth, where there are pligers but no tigers, and knows that she has not, unbeknownst to her, been transported to Twin Earth. This person seems to be in a better position to know that the singular tiger thoughts of another, who is ignorant of the facts about chemistry and Twin Earth, are indeed tiger thoughts. Here externalism seems to jettison authoritative self-knowledge, not because there is no thought of the kind available for the subject to think, but because another may be in a better position than the subject to know that the object demonstrated is indeed a tiger. (MacDonald [2006], p. 75).

The full import of MacDonald's objection may not be immediately apparent and so I will attempt to restate her objection in more vivid terms. Imagine that two subjects, Cynthia and John, are enjoying a day at the zoo. Suddenly, Cynthia notices that everything around her stops, as though frozen in time, and a genie appears before her. The genie offers to transport Cynthia and the now frozen John to Twin Earth, which is exactly like Earth except that that the striped animals in the zoos of Twin Earth are pligers rather than tigers. Though momentarily intrigued by the prospect of tricking John, who is now frozen, Cynthia ultimately declines the genie's offer to have her and John transported to Twin Earth. Then, just as suddenly as before, the genie disappears and everything unfreezes. John, who was frozen during the entire episode, is completely unaware of conversation between Cynthia and the genie. Shortly thereafter, John notices a large striped animal and pointing to it he declares, 'that tiger has stripes'. Immediately after this declaration the genie appears to both John and Cynthia. The genie explains to John that a moment earlier he appeared to Cynthia and offered to transport both her and John to Twin Earth. However, the genie does not tell John whether or not Cynthia accepted his offer. Consequently, as far as John is concerned, it remains unestablished whether the animal in the cage before him is a tiger or a pliger. I take the point of MacDonald's third objection to be that Cynthia is now in a better position than John to know that he is in fact having a tiger-type thought. For example, if the genie challenged John's claim that he was having tiger-type thoughts, it seems as though Cynthia, who knows that they are still both on Earth, is better placed to meet the genie's challenge and insist that John is in fact having tiger-type thoughts.

Notice that the present example, if successful, establishes not only that Cynthia is better placed vis-à-vis John's thought contents. Since Cynthia knows that John has not been unwittingly transported to Twin Earth, she is also better placed to know that John is having tiger-type thoughts. MacDonald summarises the upshot of her third objection as follows:

The authority which externalism purchases for another extends, not just to knowledge of which disjunct obtains, but also to knowledge of the fact, construed disjunctively, itself. And it purchases this for another precisely because that other may be better placed than the subject to know the facts about biology and Twin Earth, and to know that, unbeknownst to the subject, she has not been transported to Twin Earth. (MacDonald [2006], p. 83).

Another way of putting this point would be to say that Cynthia seems better placed (e.g., is more certain) than John even with regards to what thought he seems to be thinking.

Unfortunately, MacDonald's third objection rests on the same faulty assumption as her second—namely, that John would be unable to have tiger-type thoughts if he were transported to Twin Earth. Let us, for the sake of argument, assume the questionable premise. Since Cynthia knows that John has not be transported to Twin Earth, while John does not, Cynthia is in a better position than John to say whether or not he is having a tigertype thought. However, once we reject the questionable assumption, the conclusion no longer follows. Since John would continue to have the very tiger-type thought he is currently having, even if he were transported to Twin Earth, whether or not he has been transported to Twin Earth makes no difference vis-à-vis what type of thought he seems to be having. Of course, Cynthia is better placed than John to say whether or not the animal before him is a tiger, and is therefore able to say which of the two disjuncts obtains. However, this is consistent with McDowell's admission that it is 'possible for you to know the layout of my subjectivity better than I do in a certain respect, if you know which of those two disjuncts obtains and I do not.' (McDowell [1986], p. 154). But what Cynthia is not authoritative about is what thought John *seems* to be having, which is to say, that there is a disjunction at all.

4.5 (Re)Defining Authoritative Self-Knowledge

The notion of authoritative self-knowledge is meant to capture the idea that there is an asymmetry between our knowledge of our own thoughts and our knowledge of the external world and the thoughts of others. In §4.1 we saw that McDowell denies that infallibility is a necessary condition for authoritative self-knowledge. Thus, by his lights, the asymmetry between first and second person perspectives is not to be explained in terms of the infallibility of the former in contradistinction from the fallibility of the latter. Similarly, I do not believe that authoritative self-knowledge is best understood in terms of a subject being 'better placed' than others with regards to her thought contents. Rather, I think authoritative self-knowledge is simply meant to capture the idea that we are 'differently placed'. The 'better placed' locution is suggestive of superior access to one's mental contents while 'differently placed' merely has to do with distinctive access.

The 'differently placed' criterion for authoritative self-knowledge is summarised by McDowell:

What is special about [authoritative] self-knowledge is how it differs from knowledge of the same facts that others can have. Knowledge on the part of others has to be mediated by awareness of one's behaviour, facial expressions, and so forth. ... [A] subject is fallible as to whether her experience constitutes seeing that things are thus and so, with a fallibility that matches her fallibility as to whether things are as she thinks she sees they are. But of course that does not imply that to know that she is seeing that things are thus and so she needs to advert to the sort of 'outer' manifestations others need to advert to. (p.92)

What both the infallibility and 'better placed' conditions for authoritative selfknowledge have in common is that they presuppose that a subject has superior access to her thought contents. However, this is far from established. First, the claim that a subject has superior access to her thought contents is not true *a priori*. Moreover, most post-Freudian psychologist would view the claim with at least some suspicion, particular given the growing body of empirical evidence that suggests otherwise. Admittedly, the jury is still out on whether we have superior access to our thought contents and it is not my aim, here, to deny that we have such access. Rather, I merely wish to highlight that the fact that we have superior access to our thought contents is not something we should simply presuppose in our definition of authoritative self-knowledge.⁵⁴ We must instead be content with a much more modest claim—namely, that we have distinctive access to our thought contents.

The take-home point is this: even if another is better placed to know that I'm thinking that p in a particular instance, the fact remains that the means by which they acquired this knowledge is fundamentally different from the means by which I do. Consider some subject, S, who thinks the thought 'that tiger has stripes'. This thought is authoritative in the sense that she does not have to engage in any further exploration to know she is entertaining a tiger-thought. This is not true in the case where she is attempting to ascertain whether S^* has a certain mental content. She must rely on things like S^* 's utterance 'that tiger has stripes', reading S^* 's body language etc. Thus, the fact that S^* is not 'better placed' (a la MacDonald) than S to know that S^* is thinking that p, does not entail that S^* 's knowledge of her own thoughts is of the same character as S's knowledge of S^* 's thoughts. The fact remains that S^* 's knowledge of what S^* is thinking is so mediated. Thus, I agree with McDowell's conclusion that 'the special authority of self-knowledge consists in its not needing a basis, not in its being such that others cannot be better placed to know the facts in question.' (p. 92)

Once the above considerations are duly noted, it seems to me that that the 'better placed' objection never gets off the ground. To see this, consider the question, how does Cynthia know that John is thinking a tiger-thought? For example, suppose John is looking at a tiger and thinking a thought which he expresses with the sentence, 'that tiger has stripes'. Upon hearing this utterance Cynthia is clearly in a position to know that John is thinking 'that tiger has stripes' or (if there is no tiger available), that he is not. But notice that Cynthia only has John's utterance of 'that tiger has stripes' to go by. If John were to silently think the thought in question, then Cynthia would have no way of knowing what was going on in John's mind. It seems to me that such observations capture what is essential to our concept

⁵⁴ I believe advocates of the infallible and 'better placed' approach to authoritative self-knowledge move much too fast in that they have not even considered the possibility that while our knowledge of our inner lives may be different from that of the external world, self-knowledge may actually be inferior. After all, stranger things have been discovered to be true.

of authoritative self-knowledge.⁵⁵ On this view, John is 'better placed' than Cynthia, as far as his own thoughts are concerned, because he has unmediated access to his thoughts. He does not, like Cynthia, need to observe his own behaviour in order to know what is going on inside of his head on a given occasion. Of course, he may be mistaken about what is going on inside his head, so that he takes himself to be thinking a thought with a certain content when he does not. Thus, authoritative self-knowledge may be fallible. But this does not change the fact that when he does come to know what he is thinking, he does so in a very different way from how Cynthia comes to know what he is thinking. As far as our putative intuitions regarding self-knowledge are concerned, no stronger sense of a subject being authoritative seems necessary.

4.6 Chapter Summary

In this chapter I discussed MacDonald's argument that McDowell's singular-thought Cexternalism is incompatible with authoritative self-knowledge. I have divided MacDonald's argument into two major lines of criticism: (1) that singular-thought C-externalism conflicts with the idea that authoritative self-knowledge is infallible and (2) that singular-thought Cexternalism conflicts with the idea that authoritative self-knowledge involves that one is better placed with regards to one's own thoughts. First, I agree with MacDonald that we are not infallible regarding our singular thoughts, including singular Burgean cogito-thoughts. However, I point out that singular thought C-externalism is nevertheless consistent with the infallibility with regards to the Cartesian cogitio and with second order thoughts about how things seem to a subject. Thus, singular-thought C-externalism is still able to preserve the intuition that we are infallible regarding at least some of our second order thoughts. However, I distinguish this from the claim that self-knowledge is authoritative, which I see as rooted in the asymmetry idea. Thus, I conclude, following McDowell, that infallibility is not a necessary condition for authoritative self-knowledge.

In the second half of the chapter I employ a similar strategy in response to (2). I show that while MacDonald's arguments underscore that another may be better placed to

⁵⁵ Once the present considerations are taken into account I believe it becomes clear that MacDonald's 'better placed' criterion is irrelevant vis-à-vis privileged self-knowledge.

know if one has a certain thought content, a fact already acknowledged by McDowell, they fail to show that another can be better placed with regards to how things seem to a subject. This allows me to preserve the intuition that we are better placed than anyone else with regards to at least some of our second order thoughts, while maintaining that superior access is not a requirement for authoritative self-knowledge. I conclude that MacDonald has not given us any reason to think that singular-thought C-externalism is incompatible with authoritative self-knowledge, rightly understood.

5. THE LOTTERY ARGUMENT FOR FACTIVE JUSTIFICATION

According to McDowell, seeing that p constitutes a factive reason for believing that p. If seeing that p provides one with a factive reason for believing that p, and one's perceptual belief that p is only justified if one sees that p, then it follows that one's perceptual belief that p is only justified if one has a factive reason for p. However, the fact that the conclusion follows trivially from the two premises does not mean that the conclusion is itself trivial. Even granting the above entailment, it is possible that seeing that p provides a subject with a surplus of justification vis-à-vis a subject's perceptual beliefs. Thus, even though seeing that p provides one with factive justification, it may be that even if the entailment did not obtain (that is, if seeing were not factive), seeing that p would still be sufficient for justifying one's perceptual belief that p. To use the terminology of §3.4, even if seeing_{FACTIVE} were simply seeing (that is, seeing as conceived under the highest-common-factor conception) it may be that 'seeing' would still be sufficient to justify one's perceptual beliefs. In this chapter I will argue that only seeing_{FACTIVE} can justify.

But this claim is certainly not obvious. For one thing, my claim that justification is factive entails that a justified belief constitutes knowledge. However, it is commonly thought that a subject may have a justified belief that p even though she does not know that p. In fact, Gettier cases are typically understood to represent just such instances. Additionally, there are thought to be cases in which some proposition p is known to be sufficiently probable given some evidence {E} such that a subject who is aware of {E} would be justified in believing that p even if she recognised that she did not know that p. Take for example a subject who purchases a ticket in a fair lottery composed of 1 million tickets. On the standard⁵⁶ conception of justification a subject may be justified in believing that her ticket will lose, even though she recognises, given that the lottery is fair, that this is something she does not know.⁵⁷

⁵⁶ By 'standard' I mean the view most commonly endorsed by epistemologist, not necessarily the view that aligns most closely with our ordinary intuitions.

⁵⁷ Another example of beliefs which may be justified although they fail to constitute knowledge of the embedded proposition includes certain cases of inference to the best explanation. Consider a scientist who is confronted with two competing theories for some phenomenon. While not conclusive, the bulk of her evidence supports theory A rather than theory B. If the evidence in favour of theory A were sufficiently high, many would regard the scientist

I believe there is something right about the intuition that a subject may be justified in believing that p when she does not know that p. However, I believe there is something wrong about it as well. Below, I will attempt to distinguish between the right and the wrong. What will eventually emerge is that there is an important difference between lotterytype beliefs and perceptual beliefs, such that the latter implicates factive justification with regards to the embedded proposition while the former do not. The argument employed in the next two chapters is somewhat elaborate (convoluted?) and so a brief outline of its structure may be helpful. In the present chapter, I begin by arguing against the widely accepted thesis that a subject may justifiably believe some perceptual proposition p based on evidence that fails to guarantee that p. I do this by appealing to a lottery-type argument. A natural consequence of my view is that a belief is only justified if it constitutes knowledge, a conclusion many would find surprising. In §6 I respond to the putative objection that one may be justified in believing that p although one does not know that p. I attempt to account for the intuitions underlying this objection via a disjunctive analysis of 'I believe' statements. It should become clear by the end of the next two chapters that the notion of factive justification, with respect to our perceptual beliefs, is not nearly as counter-intuitive as it may initially seem.

5.1 Factive Justification and Truth-Guaranteeing Reasons

I wish to argue that it is a *conceptual requirement* of justification (vis-à-vis our perceptual beliefs) that it be factive. On this view, it is a conceptual requirement with regards to some reason {R}, that {R} may only justify a subject's perceptual belief that p if {R} guarantees the truth of p. When {R} meets this stipulation, I will describe {R} as a truth-guaranteeing reason for believing that p. I contrast having a truth-guaranteeing reason for p with having *evidence* for p in which evidence is seen as less than truth-guaranteeing.⁵⁸

as justified in believing theory A even if she did not know that theory A were true. Thus, there seem to be instances in which we hold that a subject is justified in believing that p when she does not know that p.

⁵⁸ Many theorists feel comfortable speaking about truth-guaranteeing evidence. My own sense is that when we think about evidence, we think about something that falls short of the fact itself. However, this question is merely terminological.

It may be tempting to contrast having a truth-guaranteeing reason for p with having defeasible evidence for p. Typically, when we describe some evidence {E} as defeasible, we mean that {E} may be evidence for p despite the fact that {E} \cup {E*} is not evidence for p. In such a case, we would say that {E*} defeats {E}, or that {E*} is a defeater for {E}.⁵⁹ However, if one has a truth-guaranteeing reason, {R}, for p, then there is no {E} that (strictly speaking) can defeat {R} since {R} already implicates truth. This final point is akin to the logical truism that one cannot make a valid argument invalid by adding premises.

However, I believe contrasting truth-guaranteeing reasons with defeasible evidence in this way is potentially misleading. For example, one may have a truth-guaranteeing reason, {R}, for *p* but fail to know that *p* because one is misled into believing that one does not have such a reason. For example, let us suppose (for the sake of argument) that seeing that it is raining provides one with a truth-guaranteeing reason for believing that it is raining. Then when Sarah looks out a nearby window and sees that it is raining this provides Sarah with a truth-guaranteeing reason for believing that it is raining. But suppose that Sarah is told that what she is seeing is actually just a large flat screen television. (Now, this is actually not the case, but we may suppose that someone decides to play a trick on Sarah.) If Sarah believes what she is told, then the report that she is looking at a flat screen television may constitute evidence that undermines her justified belief that it is raining.⁶⁰

Given the above considerations, I do not believe the contrast between truthguaranteeing reasons and defeasible evidence is very apt.⁶¹ For a more straightforward comparison I believe we must consider a subject who has evidence that only makes pprobable. For example, according to the *highest common factor* conception, (phenomenologically) indistinguishable veridical and non-veridical perceptual experiences are of exactly the same type. On this view, seeing that p cannot guarantee that p since nonveridical experiences also fall under the umbrella of seeing that p. Admittedly, most advocates of the highest common factor conception believe that seeing that p is a reliable indicator that p since most cases in which a subject sees that p, they are having a veridical

⁵⁹ See, for example, Sturgeon [1993].

⁶⁰ However, notice that (strictly speaking) it is not the case that Sarah's truth-guaranteeing reasons are defeated. Rather, the problem is that she no longer takes herself to have a truth-guaranteeing reason.

⁶¹ In this regard I part ways with McDowell who often seems to make just such a contrast. See: McDowell [1998g].

experience. Nevertheless, on the highest common factor conception, seeing that p cannot guarantee that p. It is this highest common factor conception (henceforth abbreviated as *HCF-evidentialism*) that is the primary target of the account of factive justification limned below.⁶²

On my view, having a truth-guaranteeing reason for p is disjunctively related to having HCF-evidence for p. Thus, having a truth-guaranteeing reason for p must not be equated with the conjunction of having HCF-evidence for p and the fact that p. The fundamental point being that having a truth-guaranteeing reason for p entails p. Or alternatively, having a truth guaranteeing reason is not to be confused with having nontruth-guaranteeing but undefeated evidence for p. The present contrast may be expressed in modal terms: while there are counterfactual worlds in which evidence that is undefeated in the actual world is defeated, there are no counterfactual worlds in which a truthguaranteeing reason fails to entail the truth.

5.2 Factive Justification and Lottery-type Beliefs

In defence of my position I will be appealing to a modified version of a lottery argument due to Dana Nelkin [2000].⁶³ Imagine a lottery composed of n tickets in which n is large enough to make the following claim putatively true of some particular ticket, t1, on the HCF evidential analysis: *S may justifiably believe that her ticket, t1, will lose*. For example, we can imagine a lottery composed of a million tickets in which one ticket must win but only one ticket can win. Most evidentialists would agree that *S* does not know that her ticket will lose. After all, if she did, then why would she have bothered to purchase the ticket? But according to the view I wish to impugn she may still *justifiably believe* that her ticket will lose.

⁶² The local process reliabilism discussed in §2 is an example of HCF evidentialism since it construes justification of perceptual beliefs in non-factive terms. The Gettier case against local process reliabilism fails is meant to demonstrate to inadequacy of such an approach. My diagnosis of why local process reliabilism fails is that it makes the truth-requirement external to and independent from the justification requirement for knowledge. On this score, my diagnosis echoes McDowell's: "The problem with the resources that are available in an interiorized conception of the space of reasons is that, even including the best that can be had in the way of reliability, they cannot duplicate the factiveness of epistemically satisfactory positions." (McDowell [1998h], p. 401.)

⁶³ Nelkin [2000] sets forth two separate lottery-type arguments; one involving knowledge and the other involving rational belief. My own lottery-type argument, which has to do with justification, parallels Nelkin's rational belief lottery-type argument, except that mine takes the form of a reduction while Nelkin's takes the form of a paradox.

I take it as a truism that a subject may not justifiably believe a set of inconsistent propositions which she recognises to be inconsistent. My argument will take the form of a *reductio* beginning with the assumption, '*S* may justifiably believe that her ticket, t1, will lose', and concluding with the negation of the aforementioned truism. Assuming that the first premise is the least plausible of all the premises in my argument, then my argument should establish that my first premise ought to be rejected. My reductio runs as follows:

- (A1) *S* may justifiably believe that her ticket, t1, will lose.
- (A2) If S may justifiably believe that t1 will lose, then she may also justifiably believe that t2 will lose, she may justifiably believe that t3 will lose ... she may justifiably believe that ticket tn will lose.
- (A3) *S* may justifiably believe that tickets t1, t2 ... tn will lose. [from (1) and (2)]
- (A4) *S* may justifiably believe that either t1 will not lose or t2 will not lose ... or th will not lose.
- (A5) Propositions of the following form comprise an inconsistent set: (a) p1, p2 ... pn, either not-p1 or not-p2 ... or not-pn.
- (A6) S recognises that propositions of the following form comprise an inconsistent set: (a*)t1 will lose ... tn will lose, either t1 will not lose ... or tn will not lose.
- (A7) *S* may justifiably believe a set of inconsistent propositions that she recognises to be inconsistent. [from (3), (4), (5), and (6)]

The apparent⁶⁴ far-reaching consequences of the lottery-type argument (henceforth LOTTERY) become vivid when one realises that its conclusion holds no matter how large you make n. Thus, if sound, the upshot of LOTTERY seems to be that it doesn't matter how probably one's evidence makes some proposition, so long as it fails to guarantee truth, ones belief in said proposition remains unjustified. Expressed in probabilistic terms, I take LOTTERY to show that a subject is not justified in believing that *p* based on evidence that makes *p* probable with a probability less than one.⁶⁵ More generally, I believe LOTTERY contravenes the following claim:

(A) For any subject *S*, *S* may justifiably believe that *p* based on evidence that makes *p* likely but does not guarantee that *p*.

5.3 The Justification Closure Objection

Aidan McGlynn has suggested a pair of objections to the above lottery argument.⁶⁶ First, he claims that my argument presupposes some type of closure principle. He does not actually spell out any particular closure principle, but since my argument has to do with justification he presumably has something akin to the following justification closure argument in mind:

Justification Closure Principle:

⁶⁴I say 'apparent' far-reaching consequences because whether or not the conclusion of the lottery argument may be generalised is a claim I call into question below. Ultimately, I intend to defend this claim.

⁶⁵ Note: the probability that ticket t1 will win is not a property of our knowledge about the ticket. It could not be since knowledge is factive. Rather, I attack a conception of probabilistic *justification* that gives rise to something like the following sufficiency claim:

^(*) A proposition ψ is justified if $P(\psi) > t$

Where $P(\psi)$ is a function representing the probability of ψ on the relevant evidence, and **t** is a threshold value close to one. Construed along these lines, the challenge that faces probability theorists is to characterise $P(\psi)$ in a way that is both rationally consistent and yet has a probability of less than one. The lottery argument is meant to show that no such characterisation is available. Let us grant that the P-function represents the probability that the *relevant proposition* is true, given the available evidence. In the lottery case, one's body of evidence is constituted by one's knowledge that the lottery is fair, that it is composed of a million tickets, that only one ticket will win, etc. I see this body of evidence as only making it likely that ticket t1 will lose. Now, if we see perceptual experience as presenting us with evidence that only makes *p* likely, then we are conceding that our evidence is no better than that of the lottery subject. I take the present criticism to be akin to McDowell's complaint that the interiorized conception makes truth extra (or external to) one's standing in the space of reasons.

⁶⁶ Special thanks to McGlynn for his feedback on this question.

If *S* may justifiably believe *p* and *p* entails *q*, and *S* competently deduces *q* from *p* and accepts *q* as a result of this deduction, then *S* may justifiably believe *q*.

According to McGlynn, the inference from (A3) and (A4) to the claim that what *s* may justifiably believe about this lottery has the form (a*) seems to rely on some such closure principle. However, McGlynn asks 'what mandates that [my] opponent accept such a closure principle? If anything, it looks somewhat suspect once one allows that justification does not require probability 1.'

First, I should point out that strictly speaking, the inference from (A3) and (A4) to the claim that what S may justifiably believe about the lottery takes the form of (a*) does not rest on the justification closure principle (henceforth (JCP). Rather, it rests on the following conjunction rule:

Conjunction Rule:

If *S* may justifiably believe p at time t and justifiably believe q at t, then *S* may justifiably believe p and q at t.

However, (JCP) does seem to make an appearance elsewhere in the reductio. According to the set-up of the argument, *S* is aware that one ticket must win and only one ticket can win. Given (JCP), it follows that *S* may justifiably believe that it is not the case that t1 will lose, and t2 will lose, and so on, to the millionth ticket. Given this fact, a few brief words in defence of my reliance on (JCP) may still be in order.

I begin by noting that I do not share McGlynn's intuition that my opponent has no reason (or obligation) to accept (JCP). (JCP), and closure more generally, is accepted by the vast majority of philosophers. Of course, some thinkers (most notably Nozick and Dretske) have challenged certain formulations of closure. But theirs is the minority position. (This of course does not mean that closure is true, but it seems sufficient to establish its status as the default position.) Moreover, it is worth pointing out that the version of closure implicated in the lottery argument is quite modest. If someone may justifiably believe that p, and competently deduces q from p, it seems quite reasonable to think that she would have

reasons adequate to justify her belief in *q.*⁶⁷ Interestingly, (JCP) is one of the essential premises in Gettier's original argument against JTB (See: Gettier [1963]), and rejecting (JCP) would therefore constitute a rather straightforward way resisting Gettier's conclusion. However, I think the fact that relatively few philosophers have taken this route is testimony to just how widely accepted (JCP) is and the unwillingness of philosophers in general to give it up.⁶⁸ Again, this does not establish that (JCP) is true. However, it vouches eloquently for its status as the default position in discussions of this kind.

That being said, I do not believe that LOTTERY necessarily presupposes (JCP), since *S* has independent reason for justifiably believing that t1 might not lose, t2 might not lose etc. Even if closure where false, *S*'s knowledge of how a fair lottery works, would be enough to justify her belief that any given ticket might not lose. What appears to be doing the real heavy-lifting in the lottery argument is (CR). I take (CR) to have no less intuitive support than (JCP). Thus, substantive argumentation is necessary if one is to urge the rejection of (CR).⁶⁹ Again, I believe the burden of proof rests on my opponent in this regard.

5.4 The Introspective Failure Objection

McGlynn's second objection is that even if we grant whatever closure step needed to make the aforementioned inference, premise (A6) still seems in need of defence. For example, suppose that S, due to some introspective failure, does not recognise that what she may justifiably believe about this lottery has the form of (a*). This is consistent with Ss recognizing that if she believed something which had the form (a*) she would be believing a set of inconsistent propositions. The upshot, according to McGlynn, is that (A3) and (A4) might be true, and my opponent might grant me whatever I need to conclude from that that what S may justifiably believe about this lottery is (by (A5)) inconsistent. But even in the

⁶⁷ In fact, I doubt that even Nozick and Dretske would be opposed to (JCP), since (conventionally understood) it merely implicates justified belief, rather than knowledge.

⁶⁸ See: Feldman [1995] and Audi [1995] for further discussion.

⁶⁹ Rejecting the conjunction rule regarding justification has a rather counter-intuitive consequence—namely, that a subject may be justifiably believe a set of statements she recognises to be inconsistent. While a subject who believes that each ticket will lose and also believes that not all of the tickets will lose is not believing a contradiction, it is still not possible for all the members of her set of beliefs to be true. Since a lottery subject may be aware of this fact, then rejecting (CR) means that a subject may justifiably believe a set of statements which she knows cannot all be true. Personally, such a proposal strikes me as dubious.

presence of (A6) that's not enough for (A7); it's not enough that *S* recognize that something of the form (a*) would be inconsistent – she must also recognise that what she may justifiably believe about this lottery has the form (a*). And according to McGlynn's, therein lies the rub. The number of tickets n might be very large (and the larger n is, the more plausible (A1) and (A2) are). Consequently, McGlynn questions my assumption that *S* would be able to recognize that the huge set of beliefs she may justifiably have, has the form of (a*)?

Now, I take (A6) to be true by hypothesis. Moreover, I take it to be an intelligible hypothesis, a fact that our own ability to understand the lottery argument makes obvious. But just so that it is clear that I am not begging any questions in so doing, some clarification and defence of (A6) may be in order. I do not believe that it is impossible for a lottery subject, due perhaps to some sort of introspective failure, to fail to recognise that what she may believe about the lottery takes the form of (a*). However, it think that such an oversight is no-where as likely as McGlynn makes it sound. Given the set-up of the lottery argument, the lottery subject knows all of the following:

- (i) That she is playing a lottery,
- (ii) That the lottery is composed of a million tickets,
- (iii) That one ticket must win and only one ticket can win,
- (iv) That the odds of her ticket losing are the same as any other ticket losing,
- (v) That she is no more justified in believing that her ticket will lose than any other.

Given that the subject is aware of (i)-(v), it would be odd for her not to realise that what she may justifiably believe about the lottery, given (A1), takes the form of (a*). More importantly, such an individual would most certainly be guilty of gross introspective and rational failure and may, *eo ipso*, count as doxastically irresponsible. Thus, we seem to have independent grounds for holding that such an individual is not justified.

Moreover, I think there is something misleading about McGlynn's claim that the likelihood of the LOTTERY subject failing to realise that what she may believe about the lottery takes the form (a*) increases the larger we make the overall pool of tickets. After all, I have presented the argument with a million tickets, and this fact does not make LOTTERY

any more difficult to understand than if I spoke instead of a hundred or a thousand tickets. That is to say, what is important is the formal structure of the reasoning implicated, rather than the individual application of the reasoning. *S* only needs to recognise that the inference she makes with regards to her own ticket may, by parity of reasoning, also be made regarding every other ticket, in order to recognise that what she believes about the lottery takes the form of (a*). She does not actually have to set about the arduous task of applying the relevant inference to each individual ticket. Thus, I see little motivation for thinking that a doxastically responsible subject may be mistaken on this question.

But let us grant that some given subject does not recognise that what she may justifiably believe about the lottery takes the form of (a^*). I do not see how this poses a problem for LOTTERY. Notice, the conclusion of LOTTERY merely implicates what *S* may justifiably believe, not what she actually believes. To say that one may justifiably believe *p* is to say that it is rationally permissible to believe *p*. The question with which we are presently concerned, then, is a normative one regarding what is rationally permissible for *S* to believe. The conclusion of the reductio is that it is rationally permissible for *S* to believe something she recognises to be inconsistent, and this is a conclusion I take to be unacceptable. Whether or not *S* actually does recognise that her belief is inconsistent is, as far as the present argument is concerned, beside the point.

In this regard, LOTTERY employs a similar strategy to Boghossian-type arguments against the compatibility of self-knowledge and content externalism. For the Boghossian argument to be effective, a subject need not actually come to believe some fact about her environment, for example, that she inhabits a planet with H₂O, by means of a priori introspection. In fact, given that people are not typically shuttled between Earth and Twin Earth, the scenario described by Boghossian arguments may turn out to be purely hypothetical. (Though some philosophers have argued that there are in fact real life slow-switch scenarios, such as a speaker switching between British and American English.) However, the mere fact that it is possible for the subject to acquire such knowledge is seen as sufficient to impugn the content externalist position. The intuition here is that certain epistemic achievement should not be possible. Analogously, the intuition underlying LOTTERY is that a certain epistemic achievement should not be possible.

subject believing something she recognises to be inconsistent. The point of the reductio is that we should reject (A1) since it, in conjunction with a number of other plausible premises, suggests that believing a set of propositions one knows to be inconsistent is permissible. (Again, whether or not some particular subject actually does come to believe something she recognises to be inconsistent is beside the point.)

While I believe that the McGlynn objection ultimately misses its mark, it points towards another potential objection. McGlynn challenges the move from (A6) to (A7) by pointing out that a subject might know that she is in a lottery type situation but still fail to recognise that what she may justifiably believe about the lottery has the form of (a*). But one may well ask, what about the case in which the subject fails to recognise that she is in a lottery-type case altogether. In such a case, (A6)—which LOTTERY simply assumes by hypothesis—does not even apply. LOTTERY would therefore (apparently) fail to cover such cases. This is the objection to which we now turn.

5.4 The Covert Lottery Objection:

The difficulty suggested by the introspective failure objection is remedied easily enough. However, there is a much more significant objection to the generalisability of LOTTERY, first called to my attention by Clayton Littlejohn. Littlejohn points out that the lottery argument fails to distinguish between the following two claims:

- (A) For any subject *S*, *S* may justifiably believe that *p* based on evidence that makes *p* likely but does not guarantee that *p*.
- (B) For any subject *S*, *S* may justifiably believe that *p* based on evidence that transparently only serves to make *p* likely but does not guarantee that *p*.

Littlejohn concedes that the lottery argument makes significant progress towards establishing the falsity of (B). However, he maintains that it has little or no bearing on (A). Along these lines, Littlejohn invites us to consider the example of covert lotteries, cases in
which it appears from the subject's point of view that she is not in a lottery case. Like overt lottery cases, covert lotteries implicate evidence that only makes p likely but which does not guarantee that p. However, in the case of covert lotteries the fact that her evidence only makes p likely is itself not transparent to the subject. Littlejohn summarises the upshot of the covert lotteries as follows:

It strikes me that in cases of overt lotteries, you shouldn't believe outright that claims whose truth hangs upon the outcomes of a lottery are true....However, the [lottery] argument does not apply to cases in which the truth of the relevant belief secretly depends upon how a lottery plays out.⁷⁰

Significantly, unlike the Introspective Failure objection, the Littlejohn objection does not represent an attempt to refute LOTTERY. Rather, it simply points out that LOTTERY is limited in its scope. To wit, it fails to apply to cases in which a subject does not know that her evidence is actually like that of the subject in the lottery case.

Littlejohn's distinction between overt and covert lotteries loosely parallels Jonathan Sutton's solecistic distinction between 'known unknowns' and 'unknown unknowns'. Known unknowns are cases in which the subject recognises that she does not know that p (but believes p anyway). For example, S may be fully aware that she does not know her ticket, t1, will lose and yet still believe that t1 will lose. In such a case, the claim that her ticket will lose would constitute, for S, a known unknown. I take Littlejohn's point to be something along the following lines: at best, LOTTERY only establishes that a subject may not justifiably believe an unknown. However, LOTTERY fails to establish that a subject may not justifiably believe an unknown unknown. In the case of an unknown unknown, the subject may think her evidence conclusively establishes p, and that she therefore knows p, though her evidence only makes p likely.⁷¹ Such a case would instantiate what Littlejohn calls a covert lottery. Littlejohn maintains that LOTTERY does nothing to impugn the justification of subject in the covert lottery case. However, since the evidence

 $^{^{70}}$ This was taken from an email correspondence between Littlejohn and me.

⁷¹ Ironically, I believe the lottery argument suggests just the opposite of what Littlejohn claims it does—namely, that a subject may justifiably believe a known unknown but not an unknown unknown. I will say more about this claim below.

implicated in a covert lottery is non-truth-guaranteeing, my general claim that a subject may not justifiably believe that p based on non-truth-guaranteeing evidence remains unsupported.

Before getting to my reply to the Littlejohn objection, I want to make a few important clarifications. While Littlejohn's overt vs. covert lottery distinction roughly parallels Sutton's known unknowns vs. unknown unknowns distinction, the two are not equivalent. On the contrary, I take the positions of Littlejohn and Sutton to represent two distinct alternatives to the position I wish to defend. When applied to the lottery case, Sutton's distinction already takes a stand on whether or not S knows that her ticket will lose. In both the case of known unknowns and unknown unknowns, Sutton maintains that the subject does not have the knowledge in question. Since Littlejohn concedes that the overt lottery subject does not justifiably believe that her ticket will lose, we may (assuming that justified belief is necessary for knowledge) assume that he also holds that the overt lottery subject does not know her ticket will lose. Thus, the two thinkers are on basically the same page when it comes to known unknowns.

Where the two thinkers part ways is on the question of unknown unknowns. If we conceive of the covert lottery case in terms of Sutton's unknown unknown, then we have already committed ourselves to the thesis that the lottery subject does not know that p. Littlejohn however, by choice of terminology at least, has held back from taking the additional step to assuming that the covert lottery subject does not know her ticket will lose. Moreover, once we grant that the lottery subject may justifiably believe her ticket will lose, then (assuming a JTB account of knowledge) one need only add the truth—i.e., the fact that her ticket does lose—in order to arrive at the subject knowing that her ticket will lose. In this regard, I believe Littlejohn makes room for the position which I believe the evidentialist ultimately wants to defend—namely, that one can know that p based on evidence that defeasibly supports p. (This final step becomes available once we apply the relevant intuitions elicited by the lottery case to non-lottery type propositions.) However, it is precisely this conception of knowledge (one in which truth is simply added as something extra to one's justificatory standing) that I believe makes room for Gettier-type scenarios.

Thus, my attempt to defend a JB account of perceptual knowledge that is immune to Gettier can have no part with such a conception of justified belief.

But all this is incidental to the question presently at hand. Both Littlejohn's and Sutton's distinctions allow for a subject to justifiably believe that p based on evidence that fails to guarantee that p. The fact that Sutton maintains that the lottery subject, though justified, still lacks knowledge is beside the point.⁷² By my lights, not only does the lottery subject fail to know that her ticket will lose, but she is also unjustified in believing that her ticket will lose. In this regard, both Littlejohn and Sutton represent an equal challenge to the position I wish to defend; namely that a subject may only justifiably believe that p based on truth-guaranteeing reasons for p.

It may also be instructive to restate the upshot of the Littlejohn objection in the argot of probabilistic justification. Earlier, when articulating LOTTERY, I noted that most HCF-evidentialists would grant that *S* does not know that t1 will lose. In this regard, both the HCF-evidentialist and I are on common ground. Where the HCF-evidentialist and I differ is on the question of justification. According to the HCF-evidentialist, the subject may justifiably believe that her ticket t1 will lose. Taken together, the two claims (the one on which we agree and the one on which we differ) imply that a subject may justifiably believe a proposition that she does not know. Since I deny the latter of the two claims, I am in a position to deny this implication. In fact, given my equation of knowledge with justified (strong) belief, the claim that a subject may be justified in believing a proposition she does not know seems to be one I must straightforwardly reject.⁷³ However, if the Littlejohn objection is right, LOTTERY has not yet refuted the claim that the probability theorist wishes to affirm and that I wish to deny. To wit, LOTTERY has not shown the following to be false:

(C) One may justifiably believe that *p* even if one does not know that *p*.

Rather, the lottery argument has only succeeded in refuting the stronger claim:

⁷² I have taken some creative license in construing Sutton as an opponent to factive justification in the present discussion. However, strictly speaking, my representation of Sutton's position is inaccurate since he does wish to deny that a subject may justifiably believe an unknown unknown. In employing Sutton's distinction in defence of the defeasible justification, then, I am putting it to a use which he himself does not.

⁷³ In fact, I do reject this claim, but only with a significant qualification to be spelled out below. See §6.3.2. below.

(D) One may justifiably believe that *p* even if one knows that one does not know that *p* (i.e., a known unknown).

However, as Littlejohn illustrates, the HCF-evidentialist need not be committed to (D). It bears repeating that the Littlejohn objection does not challenge the validity of LOTTERY. However, it does challenge my application of the argument. Specifically, it argues that a subject may justifiably believe based on evidence she does not recognise to be just like that possessed by the lottery subject.

5.4.1 Reply to the Covert Lottery Objection

Where the Littlejohn objection goes wrong is that it overlooks the fact that we are presently concerned with the concept of justification (understood in the abstract) rather than some particular application of the concept. Recall, according to the advocate of non-factive justification, it is part of our concept of justification that a subject's belief that p may be justified by evidence that fails to guarantee that p. However, once we admit that some type or token body of evidence does not guarantee truth, we find ourselves in the position of the aforementioned lottery subject. For example, suppose we were to hold that at best, perceptual experience provides a subject with less than truth-guaranteeing evidence for some perceptual proposition p. By adopting this posture we are establishing a paradigm of evidential justification that is akin to the overt lottery subject. Such a posture would remove all entitlement we may have for believing that p since we would be admitting that perceptual experience, at best, only makes p likely.

Interpreted along these lines, I believe LOTTERY goes some way towards establishing the falsity of the claim I took as my original target:

(A) For any subject *S*, *S* may justifiably believe that *p* based on evidence that makes *p* likely but does not guarantee that *p*. Littlejohn's distinction between overt and covert lotteries—or Sutton's distinction between 'known unknowns' and 'unknown unknowns'—is no longer salient since we are dealing with the concept of justification itself, rather than some particular application of the concept. In sum, my argument strategy may be put as follows: to suggest that the concept of justification is not factive is to suggest that evidence that fails to guarantee truth may justify. But to admit that evidence which fails to guarantee truth may justify is equivalent to saying that the (overt) lottery subject may be justified. But that the (overt) lottery subject may be justified is the very conclusion the lottery argument impugns.

Additionally, it is important to get clear on what it means for a subject to take herself to have a (truth-guaranteeing) reason, $\{R\}$, for believing p. It is not part of my view that a subject must know that some reason $\{R\}$ is truth-guaranteeing in order for $\{R\}$ to justify her belief. I do hold that for $\{R\}$ to justify a subject's belief that p, the subject must be in a position to take $\{R\}$ as her reason for her belief that p. That is to say, S must not only possess (i.e., have subjective access to) $\{R\}$, but she must also take $\{R\}$ as the reason for her belief that *p*. I also hold that {R} must be truth-guaranteeing. However, I do not hold that a subject must know (or even believe) that $\{R\}$ is truth-guaranteeing. This is based on the simple fact that one does not need to have a proper understanding of the concept of justification for one's belief to be justified. (If one did, then I fear that the beliefs of very many philosophers would be unjustified.) I suspect that it is precisely on this point that some advocates of non-factive justification go wrong. They confuse the (true by my lights) assumption that a subject may be justified even though she is not aware that her reasons are factive with the (false by my lights) claim that she may be justified by non-factive reasons. But the take-home point of the foregoing discussion is that there are two senses in which a subject may take herself to have some factive reason, {R}:

- (*) S takes herself to have {R} and {R} happens to be a factive reason for p.
- (**) *S* takes herself to have {R} and she takes {R} to be a factive reason for *p*.

My thesis clearly entails (*). However, the suggestion that (**) is required for justification entails that a subject must have a proper understanding of the concept of justification for her

belief to be justified. I want no part with such a proposal. Simply put, my view is that the concept of justification implicates factive reasons. Whether or not a particular subject happens to possess this concept is an entirely different matter.

Finally, to say that justification is factive is not to deny that a subject may be mistaken with regards to whether or not she knows that p. A subject may take herself to have a (truth-guaranteeing) reason {R}, for believing p when she in fact has no such reason. In such a case, a subject may take herself to have knowledge she actually lacks and therefore be mistaken regarding the epistemic standing of her belief that p. Thus, my position does not entail the absurd species of infallibilism according to which a subject only knows that p if she cannot be mistaken with regards to p.

5.6 Chapter Summary

This chapter represents my first step in articulating the bipartite conception of knowledge. Towards this end, I have juxtaposed truth-guaranteeing reasons with HCF-evidence. I take the lottery argument limned in this chapter to show that HCF-evidence—that is, evidence which falls short of guaranteeing truth—is insufficient for justifying our beliefs. This further motivates the disjunctive conception of experiences adumbrated earlier. If veridical and non-veridical experiences are taken to be of the same type, then perceptual experience can, at best, only make our perceptual beliefs likely. However, on the relational conception of experiences, seeing that p entails that p. Thus, seeing that p provides us with the type of truth-guaranteeing reason we need for our perceptual beliefs to be justified.

6. THE BIPARTITE ACCOUNT OF PERCEPTUAL KNOWLEDGE

My claim that justification is factive is undeniably a surprising one, not least of all because there are beliefs we take to be justified that are also false. Moreover, it has the equally surprising implication that a subject is only justified in believing that *p* if she also knows that *p*. These claims are admittedly quite unorthodox as far as the philosophical literature is concerned. However, in this chapter I hope to show that the 'bizarreness' of my claims is only apparent, and that when viewed in the right light, my view follows quite naturally from how we ordinarily think about perceptual knowledge.

6.1 Belief and the Bipartite Account

Occasionally, we use 'I believe' to identify something we are certain about. For example, I believe that I am now thinking. Sometimes we use 'I believe' to describe something we are intellectually/ rationally/ doxastically committed to. For example, I believe that 1 + 1 = 2. Often we use 'I believe' to identify something we take ourselves to know. I believe I'm currently looking at a computer monitor. This is also something that I know! However, we frequently use 'I believe' to indicate something we merely hold to be likely. Consider the following sentence:

(1) I'm not sure if it is raining, but I believe it is.

(1) seems like a perfectly natural thing to say (though on at least one reading of 'I believe' it would be self-contradictory). I take (1) to express something along the following lines:

(2) I'm not sure if it is raining, but it probably is.

Thus, often when we say that we believe that p, we simply mean that we hold p to be likely or as having a high probability of being true. However, this cannot be the notion of belief that

factors into an account of perceptual knowledge. When we come to know that p by seeing that p, we do not normally take p to be merely likely. In fact, by everyday standards, seeing that p represents the greatest amount of certainty one can have with regards to p. (This claim is suggested by clichés such as, 'seeing is believing!' or 'I saw it with my very own eyes!' and by numerous social conventions, such as the fact that in a court of law 'eye-witness' testimony is always presumed better than circumstantial evidence.) If I am right that perceptual knowledge that p requires that one be intellectually/ rationally/ doxastically committed to p being true (rather than being intellectually/ rationally/ doxastically committed to p being likely) the type of belief that factors into an account of perceptual knowledge must implicate intellectual/ rational/ doxastic commitment to the truth of p.

6.2 Belief/Knowledge Disjunctivism

There is an alternative way of making sense of the above data; namely, to see belief as disjunctively related to knowledge.⁷⁴ On the disjunctive view one either knows that p or merely believes that p. The former is factive while the latter is evidential and probabilistic, and there is no common factor shared between the two. More precisely, knowledge is not to be understood as a special kind of belief or belief plus something else (e.g., justification and truth). The disjunctive approach to belief and knowledge is suggested by sentences of the following form:

(3) 'I don't know if it is raining, but I believe it is.'

Belief/knowledge disjunctivism fits well with McDowell's conception of knowledge as a primitive. However, it flies in the face of any attempt to analyse knowledge in terms of true belief. Since I am committed to just such an analysis--namely, a JB account of knowledge—I reject a strict disjunctive conception of belief and knowledge. Before I get to my proposed account, I will outline three difficulties with the disjunctive account of belief and knowledge that further motivate my own view.

⁷⁴ The disjunctive account of belief and knowledge features prominently in Charles Travis' reading of McDowell.

Criticism 1:

Belief/knowledge disjunctivism seems out of step with much of our linguistic practice. For example, we often describe ourselves as believing things that we also take ourselves to know. I know my first name is 'Adrian', but this is also something I believe. It is difficult to reconcile the above datum with the claim that we either (exclusively) know that p or believe that p.

Criticism 2:

Belief/knowledge disjunctivism seems to present a problem for the intuition that known propositions can be justified. It is a necessary condition of a proposition being justified (as opposed to merely being justifiable) that a subject has some justification for believing the proposition. But while it makes sense to talk about someone having justification for believing that p, it does not make sense to talk about someone having justification for knowing that p. (For example, seeing that p may constitute justification for believing that p, but it does not constitute justification for knowing that p.) The latter (i.e., justification for knowing that p) would be a type of category mistake. 'Knowing' simply is not the sort of thing one could conceivably have justification for. Now, if we keep believing as a constituent of knowing (\hat{a} la the JB account of knowledge), then we may continue to hold that known propositions are justified since the believing implicated in the knowing is justified. However, once we chuck believing as a constituent of knowing, there is nothing left to be justified. Thus, we would be forced to conclude that known propositions are not justified.

Criticism 3:

Belief/knowledge disjunctivism leaves us without any means of bridging the gap between being in a position to know that p and actually knowing p. This objection is particularly pertinent to McDowell's epistemology-oriented disjunctivism (in contradistinction from, say, Snowdon's perception-oriented disjunctivism). According to McDowell's epistemic disjunctivism, perceptual experiences (i.e., the good case) put a subject, S, in the position to know that p, but it does not actually confer knowledge on $S^{.75}$ McDowell stresses that it is important to keep the two separate since there may be situations in which a subject is in a position to know that p, but still fails to know that p. For example, imagine a subject with perfectly reliable vision, who (for some unknown reason) took her vision to be unreliable. (Perhaps she was tricked into believing that her coffee was spiked by a powerful hallucinogen.) In such a case, the de facto reliability of the subject's visual experience would put her in a position to know that *p*, but her belief that her vision was unreliable would prevent her from having knowledge. Now, the question is this; what bridges the gap between a subject merely being in a position to know that p and her actually knowing that p? I submit that this is where believing (as intellectual/ rational/ doxastic commitment to the truth of p) comes in. It is only when we adopt an attitude of belief towards p that we cross over from merely being in a position to know that p to actually knowing p. Absent the attitude of belief, we may continue to have p among our mental contents (and thus be in a position to know that p) but we would not know that p. In order to know that p, one must take a stand on p; one must be intellectually/ rationally/ doxastically committed to the truth of *p*. Thus, believing *p* is indispensable for knowing *p*.

6.3 Strong/Weak Belief Disjunctivism

As an alternative to the knowing/believing disjunction I propose that a distinction be made within the concept of believing itself, what I will refer to as *strongly believe* and *weakly believe*. To strongly believe that p entails that a subject is intellectually/ rationally/ doxastically committed to p being true. To weakly believe that p entails that a subject is intellectually/ rationally/ rationally/ doxastically committed to p being likely or probable.⁷⁶ My claim that weakly believing that p takes as its object the proposition that p is likely derives from our earlier observation that statements like:

⁷⁵ See: McDowell's "Criteria, Defeasibility and Knowledge" fn 37.

⁷⁶ The primary goal of my distinction between strongly believing and weakly believing is to make sense of the different uses of "believe" in our ordinary speech. However, it is also an interesting question whether the proposed distinction may have implications for problems related to the concept of justification as well. This is a possibility I hope to explore in greater detail in the near future. But for the time being, the central contention of this post may be summarised as follows: To the extent that we are concerned with perceptual knowledge, it is belief as intellectual commitment that is at play. Thus, for a subject to know that p she must also strongly believe (i.e., be intellectually committed to the truth of) p.

(1) I'm not sure if it is raining, but I believe it is.

Are best interpreted along the lines of,

(2) I'm not sure if it is raining, but it probably is.

Now, it is a truism that some of our beliefs are more certain than others. On my view, this fact is not to be explained in terms of variations in the level of justification (Goldman's weak vs. strong justification comes to mind) but in terms of variations in the object of justification (i.e., the content of the belief being justified).77 Expressed in probabilistic terms, the object of justification for weak beliefs falls on a continuum beginning with a proposition with an estimated probability just over 0.5 and ending with propositions with an estimated probability just short of 1. Strong beliefs implicate proposition whose truth is taken to have a probability of one. More naturally, we may say that in the case of 'strong beliefs', the belief content that constitutes the object of justification is the fact that *p* simpliciter, while in the case of 'weak beliefs', the belief content that constitutes the object of justification is the fact that pis likely. Note, when I say that the object of justification for weak beliefs is 'the fact that *p* is likely', I'm not suggesting that justification itself is a probabilistic notion. Rather, I'm suggesting that the object of justification (i.e., the belief) has as its content, the claim that *p* is likely.⁷⁸ Put differently, we may say that strong and weak beliefs differ in their truthconditions. Strong beliefs are made true by the fact that p is true, while weak beliefs are made true by the fact that *p* is likely.

It is at least tempting to construe the distinction between weakly believe and strongly believe disjunctively. When S says I believe that p, this statement may be made true by S taking p to be true or S taking p to be likely. I take this disjunction to be purely semantic; rooted in our ordinary usage of the word 'believe'. We may gain a greater appreciation of the

⁷⁷ The present proposal represents a modest shift from the standard preoccupation of contemporary epistemology which has been almost exclusively concerned with justification. The suggestion is that at least some of our epistemic practice is to be explained, not in terms of varying criteria for justification, but in terms of the truth-conditions of the belief we take to be justified.

⁷⁸ Note: it may be agreed on all sides that lottery subject may justifiably believe (even know) that it is likely that her ticket will lose.

disjunctive conception of belief by comparing it with the disjunctive account of experience. For example, when the disjunctivist regarding perceptual experience (e.g., Snowdon [1988]) says that statements of the form '*S* sees O' have disjunctive truth-conditions, they mean that an utterance of such a sentence may be made true by one of two possibilities. They may be made true by *S* having an essentially relational experience implicating O, or they may be made true by *S* having an experience which is intrinsically independent of O. On the present account, something similar can be said of statements of the form 'I believe that *p*'. Such statements may be made true by a subject being intellectually/ rationally/ doxastically committed to *p* being true, or they may be made true by the subject being intellectually/ rationally/ rationa

However, we shouldn't take the parallels between the disjunctive account of experiences and the disjunctive account of belief too far. For one thing, the non-perceptual experiences which constitute the second disjunct vis-à-vis the disjunctive conception of experiences are potentially deceptive since they are indistinguishable from veridical perceptual experiences. However, there is no suggestion here that there is something deceptive about 'weakly believing' that p. Typically, when we weakly believe that p, we recognise that our attitude towards p only implicates the likelihood of p rather than knowledge of *p*. That is, weakly believing that *p* constitutes a known unknown. For an exact parallel with the disjunctive conception of experiences one must juxtapose a justified strong belief that *p* and an unjustified strong belief that *p*. That is, in the case of a subject who forms a perceptual belief based on a deceptive non-perceptual experience, she should still be regarded as having a strong belief that p, albeit an unjustified one. This goes back to the point I made earlier. When we are attempting to gain knowledge of the world by looking, it is strong belief that is implicated, since only strong belief can factor into perceptual knowledge. Thus, in the cases where one fails to have perceptual knowledge because the object of one's belief that p is false, one has an unjustified strong belief. Strong beliefs are beliefs that aspire towards (and in the good case, become constituents of) knowledge. We may sum this all up by saying that weakly believing is disjunctively related to perceptual knowledge while strongly believing is constitutively related to perceptual knowledge.

6.4 Clarifications of the Bipartite Account

This brings us to the point to which I promised to return. Earlier, I noted that my account entails the inverse of the position advocated by Littlejohn regarding the overt and covert lottery subjects. According to Littlejohn, the covert lottery subject justifiably believes that her ticket will lose, while the overt lottery subject does not. On my view the exact opposite is true. Let us begin with the overt lottery subject. Recall, for the overt lottery subject, the claim that her ticket will lose actually constitutes a known unknown. This means that the overt lottery subject only weakly believes that her ticket will lose. That is to say, the overt lottery subject is merely committed to her ticket being *likely* to lose. (Of course, given my disjunctive reading of 'I believe' statements, the overt lottery subject may still use the locution 'I believe my ticket will lose', although she actually means that she is intellectually/ rationally/ doxastically committed to her ticket being likely to lose.) Since all it takes for a weak belief that *p* to be justified is truth-guaranteeing reasons for believing that *p* is likely, the overt lottery subject may justifiably believe (understood as weakly believe) that her ticket will lose.

Although somewhat surprising, the above conclusion does not compromise, in any way, the point made by the lottery argument limned in §5.2. By my lights, the fact remains that the lottery subject, as conceived on the standard view of justified belief, does not justifiably believe that her ticket will lose. Moreover, my position offers a diagnosis of where the standard view goes wrong. According to the standard view, the overt lottery subject is justified in believing that p is true. What the lottery argument points out is that since the overt lottery subject's evidence falls short of the truth of p, she is not justified in believing that p. On my account, the lottery subject is only committed to p being likely. But the odds implicated in the lottery entails that the lottery subject's ticket is likely to lose. Thus, by my lights, the overt lottery subject's (weak) belief is justified.

In the case of the covert lotteries, I hold that the subject's belief is unjustified. This would represent a case in which a subject takes herself to have a truth-guaranteeing reason for p when she in fact does not. Interestingly, Littlejohn betrays the very intuition regarding factive justification I defend in his objection to my view. Notice, according to Littlejohn the

reason the covert lottery subject is justified is because she does not realise that her evidence is actually like that of the subject in the lottery case. On my account, what distinguishes the overt and covert lottery subject is that the former realises that her reasons are not truth-guaranteeing, while the latter does not. When the covert lottery subject's condition is re-discribed in these terms we arrive at a subject who mistakenly takes herself to have a truth-guaranteeing reason. Hence, implicit in Littlejohn's rationale for the covert lottery subject being justified is the notion that our concept of justification is truth-guaranteeing. (Or, at the very least, that the covert lottery subject takes herself to have a truth-guaranteeing reason.) Now if our concept of justification were non-truth-guaranteeing, then the covert lottery subject would be justified either way. That is, she would be justified even if she did recognise that her evidence was like that possessed by the overt lottery subject. But since Littlejohn concedes that the overt lottery subject is not justified, he has already presupposed that justification implicates truth-guaranteeing reasons. In so doing he has already given me all I need to maintain that the covert lottery subject is not justified.

6.3.1 Regarding the Covert Lottery Subject

There is an additional lesson to be learned from Littlejohn's covert lottery objection. Recall, Littlejohn claims to accept the conclusion of the lottery argument—namely, that a subject cannot be justified if she recognises that her evidence is like that of the lottery subject. That is to say, a subject is not justified if she recognises that her reason for believing that p is not truth-guaranteeing. However, Littlejohn claims that a subject may be justified in the case where she takes herself to have a truth-guaranteeing reason even though she does not. This claim may seem quite benign, until we consider its implications with regards to perceptual knowledge. Claiming that the covert lottery subject is justified amounts to the claim that a subject who takes herself to see that p when she does not may be justified. To wit, Littlejohn's position implies that a subject may be justified in believing that p even if her belief rests on a false lemma—namely, that she sees that p. However, most would regard this claim as highly implausible, a fact we would do well to make explicit in our account of justification. Thus, we arrive at the following three necessary conditions of a perceptual belief that p being justified:

- (C1) $\{R\}$ is a factive reason for believing that p
- (C2) S takes {R} as her reason for believing that p
- (C3) *S* has {R}

It bears repeating that (C1)-(C3) does not require that the subject take {R} to be factive. It only requires that she take herself to have {R} and that {R} is *de facto* factive. (C3) represents a no-false-lemma type qualification to (C2). It is not enough that one take oneself to have {R}, and that {R} happens to be factive. One must actually have {R}. If one does not, then one simply lacks the reason, and therefore the justification, that one takes oneself to have. If unexpectedly one fails to have the reason for believing that *p* one took oneself to have, then unexpectedly one fails to be justified.

There may be an appearance of conflict between my claim that the covert lottery subject is not justified and my claim that one does not need to have a proper understanding of the concept of justification to be justified. Recall, the covert lottery subject is not justified because she does not realise that her evidence is not like that of the subject in the overt lottery case. But the worry is that this appears to require that the covert lottery subject have a proper understanding of the concept of justification in order to be justified. To see that this is not the case, let us take stock of what has been said thus far. Consider the cases represented in the table below:

Subject	Has {R}	Takes Herself to Have {R}
{ S 1}	YES	YES
{ S 2}	NO	YES
{\$3}	YES	NO

Let {R} be a factive reason for p, such as *seeing that* p. Let us also stipulate that each of the above subjects believes that {R} is sufficient for her perceptual belief that p to be justified. That is, each respective subject takes {R} as her (justifying) reason to believe that p. {S1} represents the most basic case of a subject with a justified perceptual belief, while {S2} represents the most basic case of a subject with an unjustified perceptual belief. {S2} takes herself to have {R}, when in fact she does not. Consequently, I hold that {S2}'s belief that p rests on a false lemma—namely, that she sees that p. {S3} is akin to Sarah (see §5.1), who took her visual experience to be unreliable when it was in fact providing her with accurate information. Such a subject, we noted earlier, does not have a justified perceptual belief. Significantly, none of the cases described above say anything about the subject's beliefs about the nature of her justification. In this regard, they contrast with the following three cases:

Subject	Has {R}	Takes Herself to Have {R}	Takes {R} to be Truth-Guaranteeing
{ S 4}	YES	YES	YES
{S5}	NO	YES	YES
{S6}	YES	YES	NO

Once again, let us presuppose that each of the subject's take {R} to be sufficient for justifying their perceptual belief that p. Unlike the first three subjects, {S4}-{S6} do have beliefs regarding the nature of their justification. In this respect, the first three subjects may be taken to represent subjects such as infants or unreflective adults who have no definitive concept of justification. I hold that such subjects must still take themselves to have a reason for believing that p in order for their beliefs to be justified. However, they do not have to have

any particular understanding of the nature of their justification. This fact allows me to avoid an over-intellectualised conception of justification.

However, many subjects with justified perceptual beliefs do have a sophisticated understanding of the nature of their justification. {S4} represents one such subject. Not only does she have a justified perceptual belief, but she also has, what is by my lights, the correct concept of justification. {S5} matches Littlejohn's covert lottery subject. Like {S4}, {S5} has the correct understanding of what constitutes a justifying reason. (Recall my earlier observation that Littlejohn presupposes a factive concept of justification in his description of the covert lottery subject.) However, the covert lottery subject is unjustified because, just like {S2}, her belief that *p* rests on a false lemma—namely, that she sees that *p*. Notice, that {S5} is not unjustified because she has the wrong understanding of the nature of her justifying reason. On the contrary, she has what I take to be right understanding of a justifying reason. Similarly, the fact that the covert lottery subject is unjustified is not due to the fact that she has the wrong concept of justification.

The independence of a subject's justificatory status from her concept of justification is most clearly illustrated by {S6}. Here we have someone who, by my lights, clearly has the wrong concept of justification but who is nevertheless still justified. Recall, {S6} like all the other subjects, takes {R} to be sufficient to justify her belief that p. However, she also believes that {R} is a non-factive reason for p. Thus, {S6} believes that a non-factive reason may justify her beliefs that p. (In this regard, {S6} is like the defenders of the highest common factor conception of visual experiences.) Nevertheless, {S6} still meets the three necessary conditions for justification, (C1)-(C3), outlined above. (Mercifully, even those misguided philosophers who buy into the highest common factor conception may have justified perceptual beliefs.)

6.3.2 Regarding Justified False Beliefs

I believe my distinction between weakly believe and strongly believe also helps to explain another prima facie objection to the notion of truth-guaranteeing justification. The objection, simply put, is that: (E) For any subject *S*, *S* may have a justified perceptual belief that *p* even though it is false that *p*.

Now, since I maintain that justification is factive, (E) is simply false by lights. However, I believe my proposal helps to explain why (E) may seem plausible. For a weak belief to be justified it is sufficient that it is objectively probable. Such weak beliefs fall under the umbrella of known unknowns, such as (overt) lottery-type cases. It is possible for such beliefs to be justified, even while the embedded proposition, p, is false. To wit, one may have a justified weak belief that p although p is false since a false proposition may still be objectively probable or likely. This is not true, however, of our perceptual beliefs. I maintain that our perceptual beliefs are to be understood as strong beliefs. If I say that I believe that pbecause I see that p, then I am not suggesting that p is likely (as in the case of a lottery-type belief). I am taking my seeing that p to provide me with a reason to believe that p is true *tout* court. Timothy Williamson expresses a similar intuition when he observes that 'we do not usually envisage beliefs about one's own sensations as based on evidence insufficient for their truth.⁷⁹ I take this to be a distinctive feature of perceptual knowledge. I will not develop or defend this idea here. For the sake of brevity, I will simply state the relevant intuition dogmatically: I believe that it is the fact that perceptual knowledge has this feature—i.e., that it implicates strong belief and therefore factive justification—that allows it to serve as the evidential foundation for so many other types of knowledge.

6.5 Applications of the Bipartite Account

We now come to a question that may have been troubling many readers. How do any of the foregoing observations regarding lottery-type beliefs apply to perceptual beliefs? After all, our perceptual beliefs seem nothing like lottery-type beliefs. I begin my response to this question by emphasising that I do not see justification as a probabilistic notion. Hence, it is not my view that our perceptual beliefs are just like lottery-type beliefs, such as 'my ticket will lose'. In fact, that is precisely the point I have been attempting to make. The reason

⁷⁹ Timothy Williamson expresses a similar intuition when he observes that 'we do not usually envisage beliefs about one's own sensations as based on evidence insufficient for their truth.' (Williamson [2000b], p. 615).

rules like (CR) and (JCP) apply to justification but do not apply to lottery-type propositions is precisely because justification entails truth (i.e., an objective probability of one) while lottery-type propositions implicate propositions whose objective probability is less than one. In other words, my contention just is that if justification did not entail truth, we would have to treat our justified beliefs like probabilistic beliefs. But we do not treat justified beliefs in this way and this is a fact in need of explanation. My proposed explanation is that justified beliefs are factive.⁸⁰ Put slightly different: To say that justification does not guarantee the truth is to say that justification only makes our beliefs likely. But if justification only makes our beliefs likely, then (CR) and (JCP) could not apply to justified beliefs. But intuitively, (CR) and (JCP) do apply to justified beliefs. Therefore justification must guarantee truth.⁸¹ The take-home point, then, is that we are able to preserve the intuitively plausible principles of (CR) and (JCP), so long as we maintain that justification is factive.

The distinction between strongly and weakly believe also provides a solution to the so-called Preface Paradox. The Preface Paradox invites us to consider a meticulous author who takes steps to ensure that each sentence that appears in her book is true. Intuitively, it is claimed, we would consider the author justified in believing of each individual sentence that appears in her book, that it is true. However, in the preface of the book, the author concedes that there is a least one false sentence in her book. Moreover, the author seems justified in believing that there is at least one false sentence in her book. Thus, she is justified in believing a set of inconsistent propositions. Many philosophers take the Preface Paradox to show that for all its intuitive appeal, (CR) must be rejected. However, this is certainly one of those philosophical conclusions that seem acceptable because there are apparently no other options available.

However, I maintain that there is another option available. Notice that although we would regard the author's belief that each sentence is true is justified, we would not say that she knows that each sentence were true. If she did, then it would certainly be contradictory for her to both know that each sentence is true and also know that at least one of the

⁸⁰ My overall argument in this chapter may therefore be treated as an inference to the best explanation as to why (CR) and (JCP) are true of justified beliefs but not of probabilistic posits.

⁸¹ However, a qualification is now called for. I have argued that weak beliefs actually implicate the likelihood of the embedded proposition and so weak beliefs represent an exceptional case. I do not take (CR) and (JCP) to apply to weak beliefs. They had better not since that would mean that the lottery subject would not even be justified in believing that her ticket would probably lose.

sentences in her book were false. Thus, the author's beliefs that (1) every sentence is true and (2) that at least one sentence is false, are both disjunctively (rather than constitutively) related to knowledge. In other words, the author only weakly believes that each sentence is true and that at least one sentence is false. However, as we just noted (CR) does not apply to weak beliefs since weak beliefs implicate the likelihood of the embedded proposition. Thus, the fact that the author is justified in weakly believing that each sentence is correct and weakly believing that there is at least one mistake does not mean she is justified in weakly believing the conjunction of the two. This still preserves (CR) with regards to strong beliefs since strong beliefs never implicate probabilities lower than one.

Finally, we turn our attention to the perceptual Gettier case against local reliabilism described earlier, GTTRLR. In §2.4, I suggested that the diagnosis of what goes wrong in GTTRLR in terms of luck is inadequate. My alternative diagnosis of what goes wrong in GTTRLR is that Bob takes himself to have a (truth-guaranteeing) reason—namely, seeing the vase—he does not have. On this view, not only does Bob fail to know that there is a vase in the room, but he also fails to have a justified (strong) belief that there is a vase in the room.

I believe my diagnosis also helps to explain why local J-reliabilism, in both its standard and Nozickean forms, fail to provide an adequate reply to GTTRLR. The former fails, first and foremost, because it construes justification in non-factive terms. Recall, a belief forming process is reliable just in case it *tends* to produce true beliefs. Thus, according to the standard definition of process reliablism, a process does not need to always produce true beliefs in order to be considered reliable. Thus, the process implicated in the standard J-reliabilist criterion for justification is not truth-guaranteeing. The Nozickean modification to the standard process J-reliabilist picture allows us to preserve the notion of factive justification. However, it still fails because it attempts to replicate factive justification without the appeal to reasons. The fault with the Nozickean scheme, then, lies in its extreme externalism. This extreme externalism means that Nozick-style justification omits what I take to be a necessary condition for knowledge—namely, that Bob have {R} and that Bob take himself to have {R}. Moreover, it is its extreme externalism that I believe allows Nozick-style justification to be Gettiered.⁸²

⁸² The criticisms of Nozick-style justification also applies to standard process J-reliablism.

7. CONCLUSION

The account of perceptual knowledge limned in this dissertation builds on the McDowellian observation that seeing that *p* constitutes falsehood-excluding justification for believing that *p*. I see this claim as resting on four McDowellian doctrines: (1) justification internalism, (2) content externalism, (3) singular-thought disjunctivism, and (4) justificatory dogmatism. I have co-opted these theses in defence of a bipartite conception of perceptual knowledge. The bipartite account preserves the J-internalist intuition that justification is best conceived in terms of the possession of reasons. In this regard, it differs from extreme J-externalist theories such as J-reliabilism, which omit any explicit reference to reasons. However, it also seeks to preserve the J-externalist intuition that there must be an objective connection between justification and truth. On the one hand, the bipartite account agrees with the J-internalist that one's epistemic standing can only be constituted by how things are with one subjectively. But on the other hand, it rejects the interiorised view of subjectivity in favour of an externalist conception according to which objects in the external world figure among one's thought contents.

Significantly, I do not take the bipartite account to represent a definition of knowledge *simpliciter*. For example, I allow that some forms of non-perceptual knowledge may actually implicate justified weak beliefs, in which case they would not be amenable to a bipartite-type analysis.⁸³ However, this is not an issue I have taken up in the present study since I have only been concerned with a conceptual analysis of perceptual knowledge. But the question may be asked; why think that perceptual knowledge is special in this respect? While few today subscribe to the radical empiricist claim that all knowledge is rooted in human sensory experience, it is undeniable that the vast majority of our everyday knowledge is based on perceptual knowledge. Given a hierarchal conception of types of knowledge and

⁸³ A paradigm case of a type of knowledge that seems to implicate weak beliefs is knowledge of future events. I will not attempt to set forth a fully developed theory of how weak beliefs factor into knowledge about the future. But the basic idea is as follows. It is commonly recognised that human beings lack the ability to foresee the future. I may claim to know that I will be flying out of Edinburgh tomorrow, but there are numerous possible events that may prevent this from ever happening. For example, I may die in my sleep tonight, or (less dramatically) I may miss my flight due to an unexpected delay during my commute to the airport. The possibility that future events which we claim to know may never materialise is a fact we recognise even as we make such claims. I take this to suggest that our knowledge claims about the future should be seen as implicating weak beliefs (with varying estimated objective probabilities, all of which fall short of probability one). However, this is an issue best left for another time.

the foundational role perceptual knowledge plays vis-à-vis other forms of knowledge, it should hardly seem surprising that perceptual knowledge implicates a greater degree of certainty than many other types of knowledge. It is this intuition that underlies my own conviction that perceptual knowledge is best conceived on a bipartite model, which implicates factive justification.

REFERENCES

- Alston, W. [1989], *Epistemic Justification: Essays in the Theory of Knowledge*. Ithaca, NY: Cornell University Press.
- Armstrong, D. M. [1968], A Materialist Theory of the Mind. London: Routledge and Kegan Paul.
- Armstrong, D. M. [1970], 'The Nature of Mind', in C. Borst (ed.) *The Mind/Brain Identity Theory*. New York: Macmillan.
- Armstrong, D. M. [1973], *Belief, Truth, and Knowledge*. Cambridge: Cambridge University Press.
- Audi, R. [1995], 'Deductive Closure, Defeasibility and Skepticism: A Reply to Feldman', The *Philosophical Quarterly*, 45, 181: 494-499.
- Boghossian, P. [1989], 'Content and Self-Knowledge', Philosophical Topics, 27: 5-26.
- Boghossian, P. [1997], 'What the Externalist Can Know A Priori', reprinted in C. Wright et al (eds.) *Knowing Our Own Minds*. Oxford: Oxford University Press.
- BonJour, L. [1992], 'Externalism/internalism', in J. Dancy and E. Sosa (eds.), *A Companion to Epistemology*. Oxford: Blackwell.
- Brandom, R. [1995], 'Knowledge and the Social Articulation of the Space of Reasons', *Philosophy and Phenomenological Research*, 55, 4: 859-908.

Brown, J. [2004], Anti-Individualism and Knowledge. London: The MIT Press.

- Brueckner, A. [2002], 'The Consistency of Content-Externalism and Justification-Internalism', *Australasian Journal of Philosophy* vol. 80, 4: 512-515.
- Burge, T. [1979], 'Individualism and the Mental', Midwest Studies in Philosophy, 4: 73-121.
- Burge, T. [1986], 'Individualism and Psychology', The Philosophical Review, 95: 3-45.
- Burge, T. [1988], 'Individualism and Self-Knowledge,' The Journal of Philosophy, 85, 11:649-663.
- Burge, T. [1989a], 'Wherein Is Language Social?' in George, A. (ed.) *Reflections on Chomsky*. pp. 175-191. Oxford: Basil Blackwell.
- Burge, T. [1989b], 'Individuation and Causation in Psychology', *Pacific Philosophical Quarterly*, 70: 302-322.
- Burge, T. [1993], 'Content Preservation', The Philosophical Review, 102: 457-88.
- Campbell, J. [1994], Past, Space, and Self. Cambridge, MA: MIT Press.
- Chase, J. [2001], 'Is Externalism about Content Inconsistent with Internalism about Justification?' *Australasian Journal of Philosophy*, 79, 2: 227-246.
- Church, A. [1954], 'Intensional Isomorphism and Identity of Belief', *Philosophical Studies*, 5: 65-73.
- Craig, E. [1990], Knowledge and the State of Nature. Oxford: Clarendon Press.
- Evans, G. [1982], The Varieties of Reference. Oxford: Oxford University Press.

Feldman, R. [1974], 'An Alleged Defect in Gettier Counter Examples', *The Australasian Journal of Philosophy*, 52: 68-9.

Feldman, R. [1995], 'In Defence of Closure', The Philosophical Quarterly, 45, 181: 487-494.

Gettier, E. [1963], 'Is Justified True Belief Knowledge?', Analysis, 23, 6: 121-3.

Goldman, A., [1967] 'A Causal Theory of Knowing', The Journal of Philosophy, 64, 12: 357-372.

- Goldman, A. [1979], 'What is Justified Belief?' in G. Pappas (ed.) *Justification and Knowledge*. Dordrecht: D. Reidel.
- Greco, J. [2004], 'Externalism and Skepticism', in R. Shantz (ed.) *The Externalist Challenge: New Studies on Cognition and Intentionality*. New York: de Gruyter.

Harman, G., (1970], 'Knowledge, Reasons, and Causes', The Journal of Philosophy 67: 844-55.

Harman, G. [1973] Thought. Princeton, NJ: Princeton University Press.

- Harman, G. [1986] Change in View. Cambridge, MA: MIT Press.
- Hinton, J. [1973], Experiences. Oxford: Clarendon Press.
- Kripke, S. [1980], Naming and Necessity. Oxford: Blackwell.

Kvanvig, J. [1985], 'Swain on the Basing Relation', Analysis 45: 153-58.

Langton, R and D. Lewis [1998], 'Defining Intrinsic', *Philosophy and Phenomenological Research*, 58: 333-346.

Lewis, D. [1973], Counterfactuals. Oxford: Basil Blackwell.

- Mates, B. [1952], 'Synonymity', in L. Linsky (ed.) Semantics and the Philosophy of Language. Urbana: University of Illinois Press.
- MacDonald, C. [2006], 'Self-Knowledge and Inner Space', in C. MacDonald and G. MacDonald (eds.) *McDowell and His Critics*. Oxford: Blackwell Publishing.
- McDowell, J. [1986], 'Singular Thought and the Extent of Inner Space', in P. Pettit and J. McDowell (eds.) *Subject, Thought, and Context*. Oxford: Clarendon Press.

McDowell, J. [1994], Mind and World. Massachusetts: Harvard University Press.

- McDowell, J. [1998a], 'On the Sense and Reference of a Proper Name', *Meaning, Knowledge, and Reality*, 171-198, London: Harvard University Press.
- McDowell, J. [1998b], 'Truth-Value Gaps', *Meaning, Knowledge, and Reality*, 199-213, London: Harvard University Press.
- McDowell, J. [1998c], 'De Re Senses', *Meaning, Knowledge, and Reality*, 214-227, London: Harvard University Press.
- McDowell, J. [1998d], 'Intentionality De Re', *Meaning, Knowledge, and Reality*, 260-274, London: Harvard University Press.
- McDowell, J. [1998e], 'Putnam on Mind and Meaning', *Meaning, Knowledge, and Reality*, 275-291, London: Harvard University Press.
- McDowell, J. [1998f], 'Criteria, Defeasibility, and Knowledge', *Meaning, Knowledge, and Reality*, 369-94, London: Harvard University Press.

- McDowell, J. [1998g], 'Knowledge By Hearsay', *Meaning, Knowledge, and Reality*, 414-43, London: Harvard University Press.
- McDowell, J. [1998h], 'Knowledge and the Internal', *Meaning, Knowledge, and Reality*, 395-413, London: Harvard University Press.
- McDowell, J. [1998i], 'Intentionality as a Relation', The Journal of Philosophy, 95, 9: 471-491.
- McDowell, J. [2002], 'Knowledge and the Internal Revisited', *Philosophy and Phenomenological Research*, 64, 22-30.
- McDowell, J. [2006], 'Response to Cynthia MacDonald', in C. MacDonald and G. MacDonald (eds.) *McDowell and His Critics*. Oxford: Blackwell Publishing.
- Meyers, R., and Stern, K., [1973], 'Knowledge without Paradox', *The Journal of Philosophy*, 70: 147-60.
- Millar, A. [forthcoming], 'What the Disjunctivist is Right About', in A. Haddock and F. MacPherson (eds.) *Disjunctivism: Perception, Action, Knowledge*. Oxford: Oxford University Press.
- Nelkin, D. K. [2000], 'The Lottery Paradox, Knowledge, and Rationality', *The Philosophical Review*, vol. 109, 3: 373-409.
- Neta, R. [forthcoming], 'In Defense of Disjunctivism', in A. Haddock and F. MacPherson (eds.) *Disjunctivism: Perception, Action, Knowledge*. Oxford: Oxford University Press

Nozick, R., [1981], Philosophical Explanations, Cambridge: Cambridge University Press.

- Owens, J., [1990], 'Cognitive Access and Semantic Puzzles', in C. A. Anderson and J. Owens (eds.) *Propositional Attitudes: The Role of Content in Logic, Language and Mind*. Palo Alto: CSLI Press.
- Owens, J. [1992], 'Psychological Supervenience: Its Epistemological Foundation', *Synthese*, 90: 89-117.
- Peacocke, C. [1983], Sense and Content. Oxford: Oxford University Press.
- Plato [1961] Meno. Translated by Guthrie, W.K.C. In *The Collected Dialogues of Plato*, ed.Hamilton, E. and Cairns, H. pp. 353-84. New York: Bollingen Foundation.

Pollock, J. [1986], Contemporary Theories of Knowledge. Savage, MD: Rowman and Littlefield.

Pollock, J. [1987], "Defeasible Reasoning", Cognitive Science 11:481-518.

Pollock, J. [1995], Cognitive Carpentry, Cambridge, Mass.: MIT Press.

- Putnam, H. [1963], 'Brains and behavior' in R. Butler (ed.) Analytical Philosophy: Second Series.
 Blackwell. Reprinted in Mind, Language, and Reality, Cambridge: Cambridge University Press, 1975.
- Putnam, H. [1975a], 'The Meaning of "Meaning"', in K. Gunderson (ed.) Language, Mind, and Knowledge, vol 2. of Minnesota Studies in the Philosophy of Science, ed., Minneapolis: University of Minnesota Press.
- Putnam, H. [1975b], 'Explanation and Reference', *Mind, Language, and Reality*. Cambridge: Cambridge University Press.

- Putnam, H. [1975c], 'Is Semantics Possible?', *Mind, Language, and Reality*. Cambridge: Cambridge University Press.
- Russell, B. [1905], 'On Denoting', Mind, 14: 479-493.

Ryle, G. [1949], The Concept of Mind. New York: Hutchinson and Co.

Sellars, W. [1963], Science, Perception and Reality. London: Routledge.

- Snowdon, P. F. [1990], 'The Objects of Perceptual Experience', *Proceedings of the Aristotelian* Society Supplementary Volume, 64: 121–150.
- Spohn, W. [2002], 'A Brief Comparison of Pollock's Defeasible Reasoning and Ranking Functions', *Synthese* 13:39-56.
- Stalnaker, R. [1968], 'A Theory of Conditionals', *American Philosophical Quarterly*, monograph no. 2: 98-112.
- Sturgen, S. [1993], 'The Gettier Problem', Analysis, 53,3: 156-164.

Sutton, J. [2005], 'Stick to What You Know', Nous, 39, 3: 359-396.

Vahid, H. [2003], 'Content Externalism and the Internalism/Externalism Debate in Justification Theory', *European Journal of Philosophy*, 11: 89-107.

Vogel, J. [2000], 'Reliabilism Leveled', The Journal of Philosophy. 97, 11: 602-623.

Williamson, T. [2000a], Knowledge and Its Limits. Oxford: Oxford University Press.

Williamson, T. [2000b], 'Scepticism and Evidence', *Philosophy and Phenomenological Research*, 60, 3: 613-628.