

A Defence of Analyticity

Submitted by David Ian Tait

University College London

Submitted for the Doctor of Philosophy degree

I confirm that the work presented in this thesis is my own and the work of other persons is appropriately acknowledged.

Abstract

There is *prima facie* reason to suppose that there are analytic truths, our knowledge of which is explained simply by our understanding them. One recent line of argument challenges this view on the grounds that, for any given proposition, it is always possible to understand it without knowing it. If understanding is to explain our knowledge of certain truths, then, how is it possible for someone to understand them and yet fail to know them?

We can accommodate these cases of disagreement by construing the epistemic state in which a subject is placed by understanding an analytic truth as one of *being in a position to know*. In understanding an analytic truth, a subject may have the epistemic resources required for knowledge and yet be unable to exploit this position; this allows for the possibility that in those cases where a subject does know such a truth, the knowledge is explained by the subject's understanding. This sense of *being in a position to know* receives support from the need for such a notion in describing certain features of our perceptual knowledge.

Understanding an analytic truth enables a subject to recognise that its truth-conditions must be fulfilled. This is ultimately made possible by there being certain propositions that have the status of structuring the linguistic practice in which the subject participates. These propositions are held fixed as we evaluate the possible ways that the world could be and so come out as true in all possible worlds. A subject who is sufficiently integrated within the practice and who understands an analytic truth is thereby in a position to recognise its status within the practice. Using this model we can identify two kinds of disagreement consistent with the claim that understanding an analytic truth puts one in a position to know it.

Acknowledgements

It seems appropriate that the culmination of such a project should acknowledge those whose help has, in various ways, contributed to the process of its development. I am grateful to Rory Madden who, while supervising me as I wrestled with the ideas for a completely different topic, suggested to me that my writings about analyticity harboured the prospects of a much more fruitful thesis than I was then on course to produce. I suspect his judgement was right, and I am thankful he chose to share it. I am even more indebted to those under whose guidance the kernels of those initial ideas were developed: Mike Martin, whose understanding of ‘the big picture’ against which the topic can be seen to take shape has been invaluable, and Marcus Giaquinto, whose concern for detail has kept me from numerous mistakes and pushed me into thinking much more precisely about the issues here.

A central part of the thesis was presented at the MindGrad conference at the University of Warwick in December 2009. I would like to acknowledge my gratitude for stimulating and instructive comments from Bill Brewer, Quassim Cassam, Tim Crane, Guy Longworth, and in particular my respondent in that session, Matt Soteriou.

Thanks is owed to my friend Matthew Manning for listening patiently through the development of almost every idea – and many others besides – contained herein, as well as for his sensible advice. Thanks also to Craig French for kindly agreeing to read through the work in its final stages and demonstrating such a keen editorial eye. His help made the last few days of the process much more manageable.

I want to express my deep gratitude to my parents for their unwavering support and encouragement which has meant so much to me and without which the project would not have been possible.

Finally, I must thank my wife, Cheryl, for making things better when the work wasn’t going so well and for making them even better when it was.

Table of Contents

Introduction	6
1. Truth in Virtue of Meaning	19
2. The Epistemological Conception and Williamson's Challenge	52
3. Perceptual Knowledge and Being in a Position to Know	69
4. Analyticity, Understanding and Being in a Position to Know	93
5. How Could Understanding Explain Our Knowledge of Analytic Truths?	113
Conclusion	150
Bibliography	155

Introduction

(i) Introducing analyticity

Let us start with some basic data. There is an intuitive difference between the way in which we know that vixens have dichromatic vision and the way in which we know that all vixens are female foxes. In trying to articulate this difference, we are inclined to say something along the following lines: whereas our knowledge in the former instance derives from investigation into how things are in the world, our knowledge in the latter instance does not seem to require investigation of this sort. Rather, so this intuition goes, our knowledge seems to derive merely from our understanding of the words involved, or if you prefer, from our grasp of the proposition that they express.

These initial intuitions suggest a peculiar epistemological phenomenon. For insofar as we must understand any claim in order to form a judgement about its truth, understanding is always a feature of our knowledge. But in typical cases, the role that understanding plays is limited to our grasping some proposition, the warrant for which is acquired in experience. However, in the case of ‘All vixens are female foxes’, we have the intuition that once we grasp the proposition, there is no warrant from experience needed in order to know it. It seems that understanding alone provides us with the resources necessary for knowledge.

Kant, in the *Critique of Pure Reason*, discusses the phenomenon whereby our understanding has the capacity to furnish us with knowledge, labelling ‘analytic’ those truths capable of being known in this way. His discussion of analyticity focuses on statements of a subject-predicate form, leading him to characterise an analytic truth as one in which “the predicate B belongs to the subject A as something that is (covertly) contained within this concept A.” (Kant 1998, A6) Put this way, Kant’s remarks have the appearance of being more interested in a linguistic, or conceptual, phenomenon concerning the relations between words or concepts in a way that might seem more familiar to modern lexical semantics. However, the epistemic import of the characterisation is spelled out when he writes that an analytic truth is one in which “I do not need to go beyond my concept at all in order to formulate the judgment, and

therefore need no testimony from experience.” (1998, B11) Since it is this epistemological phenomenon we are really interested in, we may dispense with Kant’s overly restrictive “containment” metaphor and extend the notion of analyticity so as to include all instances where our grasp of a proposition provides us with the resources by which we know it. Once extended in this way, we leave open the possibility that analytic explanations may be found for the whole range of a priori truths, including the truths of logic, mathematics and a broad range of conceptual truths (such as ‘No object can be both red and green all over at the same time’).

Despite the intuitive appeal of supposing that understanding explains our knowledge of at least some a priori truths, the notion of analyticity has suffered something of a chequered history over the past century or so. Here is what I take to be the standard story that receives widespread acceptance among many contemporary philosophers. The logical positivists, trying to find an explanation for our a priori knowledge that would avoid the kind of metaphysical excesses they eschewed, appealed to the idea of analyticity as something that offered “truth in virtue of meaning”. In this way, they were able to maintain that there is nothing metaphysically odd about such knowledge, since it is really just knowledge of linguistic conventions and thus not really about the world in any substantive way. Quine, however, demonstrated this stance to be untenable, arguing that a proposition cannot be regarded as true in virtue of meaning alone but rather can be true only because the world is as the proposition represents it to be. Since all propositions are made true by how things are in the world, our acceptance of any proposition is always subject to revision according to its conformity to the best account of our experience. There can, therefore, be no propositions with the privileged status of immunity from revision that would come with truth in virtue of meaning alone. Following Quine’s attack on analyticity, the notion fell into disrepute and lay moribund until it was offered the chance of rehabilitation by Boghossian. Boghossian distinguished the metaphysical issues concerning that in virtue of which a statement is true from the epistemological issues concerning how it might be known. In drawing this distinction, Boghossian separated the dubious positivistic aspects of analyticity from the more valuable idea that there be a class of truths our understanding of which is sufficient for us to know them. This opened up

the possibility of an account that could evade the Quinean objections while delivering all that we really wanted from the notion of analyticity in the first place.

I call this the standard story because I think that something very much like the foregoing narrative has come to shape much of the recent literature. Even amongst those who have subsequently questioned Boghossian's strategy of separating the metaphysical and epistemological conceptions, or who have expressed doubts over whether the strategy yields a defensible account of analyticity, there is a fairly prevalent consensus concerning the contours of the story.¹

We might add, as a sort of coda to the standard story, the more recent development of Williamson's sustained criticism of the epistemological conception of analyticity. For while Williamson takes the view that Boghossian's proposal affords the notion a cogency that eluded the metaphysical conception, he has argued repeatedly that there is no viable account of understanding that is capable of underwriting or supporting analytic explanations of our knowledge. Since it is always possible for a subject to understand a putatively analytic truth without knowing it, Williamson argues, there can be no truths our knowledge of which is explained simply by appeal to our understanding alone.

If Williamson is right then, contrary to our initial intuitions, there are no analytic truths. Yet these intuitions remain fairly strong. This, to my mind at least, makes this recent development in the debate quite fascinating and worthy of further exploration. Consequently, a relatively large portion of this thesis is taken up with developing a defence against Williamson's objections. However, the account that emerges is one that ends up dissenting from the main themes of the standard story. Once we accommodate the insights within Williamson's argument and develop a view of how understanding is capable of explaining our knowledge of analytic truths, we see that certain distinctive elements of the metaphysical conception resurface, giving lie to the commonplaces that Quine decisively refuted the views of those who defended this conception and that Boghossian succeeded in isolating two entirely distinct conceptions of analyticity.

¹ See, for example Margolis and Laurence (2001), Williamson (2007), Russell (2008) and Garcia-Carpintero and Otero (2009).

Whatever view one takes concerning the brief narrative that has been sketched, which is amplified and precisified over the next two chapters, it at least serves to delineate a core set of issues to be addressed in exploring the notion of analyticity. First, if we take seriously the initial intuitions that suggest there may be some truths our knowledge of which is explained by our understanding them, then we require some account of how understanding could be responsible for any such achievement. This is the key concern for any defence of analyticity. The challenge lies in the need to motivate an independently plausible conception of understanding that is capable of doing the required epistemic work – which, we will see, is no easy task. Second, we might like some explanation for the common intuition that there is something slightly peculiar about counting some of the paradigmatic candidates for analyticity – particularly the conceptual truths – as items among our knowledge about the world when many of them have the flavour of being rather truistic. It is surely relatively easy to sympathise with those who, when first introduced to the notion of analyticity through classic instances such as ‘All bachelors are unmarried males’ and ‘All vixens are female foxes’, insist that these are not really features of our knowledge about the world but, rather, mere definitions of the relevant expressions. It is this thought, I take it, that perpetuates the appeal of characterising analyticity as “truth in virtue of meaning” despite the various objections raised against the phrase. Even if the defender of analyticity were ultimately to reject the positivist views associated with that slogan, we might still regard it as desirable that an account should offer some sort of explanation for the prevalence of the underlying intuition. Third, in light of Williamson’s recent objections against analyticity, we need some story about how it is possible for a subject to understand an analytic truth and yet fail to know it.

Though clearly separable, these issues need not be viewed as independent of one another. According to the standard story outlined above, the logical positivists can be understood as having allowed their preferred answer to the second issue to shape their response to the first. On the other hand, we might regard someone like Boghossian as allowing his dissatisfaction with the positivist approach to focus his efforts on giving primacy to the first issue, only to then have difficulty in answering the third. Each issue impacts on the others. What we desire is an account capable of offering a unified response to all three.

At the risk of appearing contrary, since it is hardly the obvious point at which to start, my primary focus here is on the possibility of accommodating cases of disagreement, in which a subject understands an analytic truth and yet fails to know it. This is partly because it is perhaps the most topical of issues relating to analyticity in recent literature, with several philosophers – Williamson most notable among them – arguing that it presents a serious difficulty for analytic explanations of the a priori. It is also partly a reflection of what is at stake in the debate; if Williamson’s criticisms were right, and the consequence is that there can be no analytic truths, then the other two issues would stand in considerably less need of address. A third reason to focus on this issue is that once we recognise the truth in Williamson’s point about the need to accommodate the possibility that someone could understand an analytic truth without knowing, we will need an account of how understanding could explain our knowledge that is sensitive to this requirement.

(ii) Analyticity and rational intuition

Historically, analyticity has been of interest as a means of explaining our a priori knowledge without having to postulate other, perhaps more metaphysically extravagant, modes of explanation. In particular, the traditional rivals to analytic explanations of our knowledge are those which appeal to some sort of faculty of rational intuition, by which we ‘see’ or intuit the truth of certain propositions. We might regard it, then, as a dialectical constraint upon any attempt to defend analyticity that it should not ultimately end up collapsing into an account of rational intuition. The problem lies in knowing where to draw the line between them. Each type of explanation has its greater and lesser extremes, but upon examination of their more moderate versions, the boundary separating them into their respective sides becomes less sharp. We have said, in introducing the notion of analyticity, that it is the phenomenon whereby understanding or grasp of a proposition explains our knowledge of it. Analyticity, then, is an understanding-based account of our a priori knowledge. Yet some defenders of rational intuition maintain that we can construe this faculty as nothing more than our ability to understand and reflect upon propositions of a certain kind,² in which case rational intuition may also be categorised as an understanding-based account the a priori, thus closing the gap between the two types of explanation.

² BonJour 1998: 109

Of course, we might wonder whether there is a genuine dispute to be had here; it is possible that if we flesh out the details of a moderate account of rational intuition and a moderate account of analyticity that we might discern little difference between them, instead finding that they are in fact much closer to one another than either happens to be with the extreme versions of their respective sides in the debate. But even if this is the case, we might nevertheless look for some criterion that would at least allow us to distinguish a paradigmatic account of analyticity from one of rational intuition.

Part of the difficulty, some defenders of analyticity have complained, is that it is not especially clear what rational intuition is supposed to be; Boghossian writes,

The central impetus behind the analytic explanation of the a priori is the desire to explain the possibility of a priori knowledge without having to postulate such a special faculty [as rational intuition], *one that has never been described in satisfactory terms*. (Boghossian 1997: 334, italics added)

Elsewhere he says,

The single most influential consideration against rational insight theories can be stated quite simply: no one has been able to explain – clearly enough – in what an act of rational insight could intelligibly consist. (2001: 635)

At one end of the spectrum, we have the Gödelian idea of a quasi-perceptual faculty that offers direct insight into meanings or concepts, and on which we ‘see’ an a priori proposition to be true.³ At the other, we have the insistence that all there is to our knowledge of a priori propositions is the ability to understand them and recognise that they must be true.⁴ With the former, we might wonder how we happen to be related to meanings or concepts in the relevant way, or even what reason we have to suppose that we are. With the latter, we might wonder how it is that understanding by itself could furnish us with the alleged ability to recognise that a given proposition must be true. Although we may grant that the same question needs addressing by the defender of

³ Cf. Peacocke 2000: 261-2

⁴ BonJour 1998: 101, 109

analyticity, it is fair to say that answers to these questions have been conspicuously lacking from those who align themselves with the idea of rational intuition, a point acknowledged by one of its principal contemporary advocates. Responding to a demand that a satisfactory defence of rational intuition should “somehow lay bare, in appropriate detail, how some capacity that we have gets to work on the properties we are able to think about so as to disclose their nature,” BonJour writes that “neither my rationalist view nor any other that I am familiar with has even come close to providing the sort of explanation that this objection demands.” (BonJour 2001: 673)⁵

BonJour points out that the label of ‘rational intuition’ offers some clue as to what its proponents have supposed it to be. The point of referring to the source of the knowledge as ‘intuition’ underlines the idea that the knowledge is not derived from any inferential process or through any axioms or principles already accepted. Rather, it is direct and immediate. Calling it ‘rational’ highlights the commitment by its advocates to the idea that such a faculty is “intellectual or reason-governed, anything but arbitrary or brute in character.” (BonJour 1998: 102) We see this in BonJour’s description of the process by which knowledge is acquired in this way:

when I carefully and reflectively consider the proposition (or inference) in question, I am able simply to see or grasp or apprehend that the proposition is necessary, that it must be true in any possible world or situation.... (BonJour 1998: 106)

There is, on this view, nothing more to be done than to consider the proposition. The rationality of the process is suggested by the talk of ‘seeing’, which with its perceptual overtones connotes the thought that there is something whose truth we are tracking by this method, our responsiveness to which is explanatory of our knowledge. However, it is not clear that these considerations are going to distinguish an account of rational intuition from one of analyticity. For the defender of analyticity may equally insist that in all relevant cases our knowledge of analytic truths is unmediated and direct in the same way, acquired only through a subject’s understanding or grasp of a proposition. It is, of course, unsurprising that there should be at least this much similarity between the explanations,

⁵ Perhaps somewhat surprisingly, he goes on to add, “I do not think that it follows in any clear way that the idea of rational insight should be rejected.”

since they are both trying to account for the same phenomenon. But it means that characterising the explanations in terms of how things seem to the subject is only going to leave us wondering where to draw the distinction between them.

A more promising avenue is to consider the use to which the respective explanations have been put by their defenders.⁶ For whether or not we any longer suppose the labels ‘rationalist’ and ‘empiricist’ to have meaningful application in contemporary philosophy, it is undeniable that the historical use of these terms has tended to reflect two different views about a priori knowledge. Moreover, these two different views have corresponded to the division between those explanations that appeal to rational intuition and those that appeal to analyticity.⁷ Given this, if we identify the attitudes to the a priori that motivated rationalists to appeal to rational intuition or empiricists to embrace analyticity, we might begin to have some criterion by which we can distinguish their accounts.

Writing about the attitudes that characterise these positions, Ayer says that a rationalist holds the view that “there are some truths about the world which we can know independently of experience,” (Ayer 2001: 66) whereas an empiricist denies that we can have a priori knowledge of reality. If this is right, we could divide rationalists and empiricists by their stance with respect to (P),

(P) A priori knowledge of reality is possible.

But this raises a further question concerning what should count as a priori knowledge of reality. We might suppose that our knowledge should count as being ‘of reality’ if the propositions in question are, in some sense, about the world and are true by virtue of how things are in the world. This is certainly something that some empiricists, including Ayer and Carnap, have denied in relation to analytic truths. Their reasons for thinking along these lines are discussed in greater detail in the next chapter. In brief, however, they take the view that an analytic truth can be known without investigating the world only

⁶ This portion is influenced in significant respects by a very helpful discussion of the surrounding issues in Cassam (2000).

⁷ I am ignoring, for present purposes, the more radical empiricism that denies the possibility of a priori knowledge entirely.

because it is not really about the world. The proposition that either some ants are parasitic or none are, Ayer says, “provides no information whatsoever about the behaviour of ants, or indeed, about any matter of fact.” (Ayer 2001: 73) Analytic truths thus do not supply us with any new information about how things are in the world, but rather “call attention to linguistic usages.” In this way, they lack factual content and are not to be regarded as purporting to be about the world, in the relevant sense.

Ayer’s criterion, however, does not afford us a clean separation of the rationalist and empiricist views. For there are empiricists who advocate analytic explanations of the a priori but equally insist that analytic truths are both about the world and made true by how things are in the world. Boghossian, endorsing Quine’s criticism of the positivists on this point, asks,

What could it possibly mean to say that the truth of a statement is fixed exclusively by its meaning and not by the facts? Isn’t it in general true – indeed, isn’t it in general a truism – that for any statement S,

S is true iff for some p, S means that p and p? (Boghossian 1997: 335)

And again he writes,

In general, I have no idea what would constitute a better answer to the question: What is responsible for generating the truth of a given class of statements? than something bland like ‘the world’ or ‘the facts’. (Boghossian 1997: 336)

So we cannot invoke the issue of the factuality of a priori knowledge in a way that will characterise the empiricist view in its most general form. BonJour seems to get closer to an acceptable drawing of the distinction when he observes that rationalists have typically been impressed by the idea that the mind intuitively grasps or apprehends necessary facts about the structure of reality,⁸ whereas empiricists have wanted to maintain that this knowledge is restricted to “propositions that reflect relations among our concepts or meanings or linguistic conventions, rather than to those that make substantive claims

⁸ BonJour 1998: 15-16

about the character of the extra-conceptual world.” (BonJour 1998: 18) This certainly captures the initial appeal of the rationalist view that our a priori knowledge extends to substantive truths about the world and is not “limited in its scope to tautologies or matters of definition.” (1998: xi) And it accurately represents the empiricist stance that the explanation for our a priori knowledge is ultimately rooted in the meanings of our words or the propositions that these express. But BonJour’s characterisation of the division is misleading to the extent that it suggests that what is at stake between the rationalist and the empiricist is their respective attitudes concerning the substantiveness of our a priori knowledge. For it is possible to advocate an understanding-based, analytic explanation of the a priori while maintaining the truths that are known in this way are nevertheless substantive truths about the world. Boghossian, for example, argues that to possess the concept *if* is to be entitled to endorse instances of modus ponens, thereby offering a broadly analytic explanation of our knowledge of a certain inference pattern.⁹ But there is no suggestion that this knowledge is in any way insubstantial; Boghossian thinks it a fact about the world that the pattern in question is one in accordance with which we ought to reason and as such is surely to be counted as substantial.¹⁰

Ayer’s and BonJour’s efforts to characterise the divide between rationalism and empiricism illustrate just how tricky it can be to identify what it is that separates them. And we might have to accept that there is no one way of characterising the debate that is not subject to amendment or qualification for some peripheral view. However, I want to suggest that we can locate the central difference between the views in what they take our a priori knowledge to be tracking, that is, in what they suppose our a priori knowledge to be responsive to. Our perceptual knowledge is the product of tracking, or being responsive to, the objects in our environment; we form our perceptual beliefs in response to the information that is presented to us by our perceptual faculties concerning how things are in the world around us. In these cases we are responsive to things clearly outside of ourselves. Our introspective knowledge of our beliefs, intentions and desires is the product of our ability to track our own mental states; we form these beliefs in

⁹ Boghossian (2003a)

¹⁰ Carnap too seems to suggest that it does not follow from a truth’s being analytic that is therefore insubstantive when he writes that the statement “five is a number” “is rather trivial (in contradistinction to a statement like “There is a prime number greater than a million”, which is likewise analytic but far from trivial).” (Carnap 1956b: 209)

response to what we are able to discern about ourselves by this method. In these cases we are responsive to things internal to ourselves. Similarly, when we seek an explanation for our a priori knowledge, we may regard ourselves as seeking an answer to the question: to what is our a priori knowledge responsive? And we can view the division between the rationalist and the empiricist as emanating in their responses to this question.

The rationalist, we have already observed from BonJour, is drawn to the idea that our a priori knowledge constitutes substantive knowledge about the structure of reality. The appeal to a faculty of rational intuition is deployed to capture this idea that we have some sort of direct insight into certain truths about the world. On this view, our a priori knowledge is the product of tracking, or being responsive to, things – perhaps Platonic entities – that are external to us.¹¹ A rationalist may insist that theirs is an understanding-based approach to the a priori insofar as the method for intuiting an a priori truth might involve a subject doing nothing more than understanding and reflecting upon the relevant proposition, but that in doing so the subject is nevertheless responsive to things external to himself. On the other hand, it would seem to be characteristic of the empiricist's view that our a priori knowledge is not a direct and immediate insight into the structure of reality but rather born of our concepts and the meanings of our words. This, as we shall see later on, does not commit us to the claim that the resultant knowledge is not in any way about the world; but we might take the view that it is not the world to which we are *directly* responsive when we know an a priori truth. To take the relatively simple case of 'All vixens are female foxes', an empiricist can accept that this is in some very straightforward way about vixens and thus about the world. But even though the world is its subject matter, the non-experiential nature of the knowledge suggests that we are not responsive to the world in the same way as when we see a vixen and, on the basis of our perceptual experience, note that it is brown. Rather, the empiricist maintains, we are responsive to considerations concerning our concept possession or our linguistic practice. In this way, we track something which, if not internal to ourselves (if meanings and languages are social, externally determined things), is certainly not anything over and above what we already have at our disposal as language users.

¹¹ Cf. BonJour 1998: 156-161

This way of drawing the distinction between rationalism and empiricism affords a means of characterising the difference between appeals to rational intuition and analytic explanations of our a priori knowledge. An account of rational intuition is one that takes our a priori knowledge to be directly responsive to objective features of the world which are external to us and our linguistic practices; an analytic explanation of our knowledge supposes it to be tracking some feature of our linguistic or conceptual practice. It captures the fairly intuitive divide between the two views in a way that preserves the classifications most major players in the literature have assigned to themselves: Ayer, Carnap and Boghossian in the empiricist camp; Bonjour in that of the rationalist. (One notable exception here is Peacocke, who describes himself as a rationalist, though much of his work concerns the knowledge that arises from our possession of certain concepts, which would by these lights place him as an empiricist; I do not find this too troubling a result, being inclined to regard him as championing the possibility of substantive a priori knowledge *per se*, rather than as a partisan advocate of one side or other in the manner of someone like Ayer or Bonjour.)¹² On this basis, although the distinction remains somewhat rough-and-ready, and no doubt subject to further borderline cases, I will proceed on the assumption that something along these lines captures the division between accounts of rational intuition and analyticity.

(iii) The structure of the thesis

The central claim of this project is that, despite recent objections to the contrary, it is consistent with an analytic explanation of the a priori that there can be cases of disagreement in which a subject understands an analytic truth but fails to know it. Two questions are immediately raised by this. First, if understanding is to explain our knowledge of analytic truths in such a way as for it to be possible that someone might understand an analytic truth and not know it, what is the epistemic state in which a subject is placed by understanding an analytic truth? The answer, in short, is that understanding an analytic truth puts one in a position to know it, where this notion of *being in a position to know* accommodates the kind of defeasibility that is supposed to make disagreement so problematic for an account of analyticity. Second, following this,

¹² In addition to this, in conversation Professor Peacocke has been content to describe the account he has defended in his work as a form of analyticity.

how does understanding an analytic truth put one in a position to know it? I argue that the explanation is ultimately rooted in a subject's participation in the practice of a linguistic community, in which the analytic truths are those held fixed in such a way as to structure that practice; a subject's participation enables him to recognise the status these propositions have in the practice of the community.

In chapter one I explore the idea of analyticity as "truth in virtue of meaning," with particular reference to the development of this strategy by Carnap, one of its foremost advocates. I argue that contrary to the claims of those who subscribe to the standard story, this conception of analyticity is more robust than is often supposed; presenting the famous, protracted disagreement with Quine, I argue that neither it nor other objections commonly offered against it tell decisively against it. In chapter two I turn to the epistemological conception favoured by more recent defenders of analyticity who have wanted to avoid the Carnapian strategy. Proponents of this conception have typically sought to establish that knowing an analytic truth is constitutive of understanding it. The problem for this approach is that it is possible to understand an analytic truth and not know it. I argue that we can accommodate this possibility by construing the epistemic state in which a subject is placed by virtue of understanding an analytic truth as one of *being in a position to know*. In chapter three I show that the appeal to being in a position to know is supported by the need for the same notion in describing certain features of our perceptual knowledge.

Williamson claims that a conception of understanding sufficiently weak to accommodate cases of disagreement will be too weak to underwrite an account of analyticity. I examine his argument for this view in chapter four. I argue that we can accept the broad features of the picture he offers of our linguistic practice without having to accept the weak conception of understanding he supposes to follow from this. In chapter five I show how we can use the emergent picture to explain how understanding an analytic truth could put a subject in a position to know it.

Truth In Virtue of Meaning

Insofar as the sentences of mathematics refer to reality, they are not certain, and insofar as they are certain, they do not refer to reality. (Albert Einstein, 'On the Electrodynamics of Moving Bodies')

We typically regard it a norm governing our acceptance of declarative sentences, or the propositions these express, that we should accept them if and only if the world is as they represent it to be. The idea, on this standard view, is that the meaning of a sentence demarcates some way for the world to be and is true if and only if the world happens to be that way. The world thus has an essential role in what it is for something to be true. If we discover the world to be some particular way, we have an excellent reason to accept a proposition that represents it to be so. If, on the other hand, we were to learn that there was nothing about the world that corresponded with the way it is represented to be by a proposition, we would – I presume – take this to be grounds for withholding acceptance of the proposition in question.

This characterisation of our practice draws no distinctions between different types or classes of declarative sentence to which it applies especially, for the simple reason that it is a general claim that we suppose holds for all such sentences equally. There is an array of sentences that we hold as expressing truths: 'Grass is green', 'Cambridge is north of London', 'There is a prime number greater than twelve', 'All vixens are female foxes', 'Everything is self-identical'. Yet in all of these cases we may feel that we are justified in holding them true precisely because they each represent the world as being some particular way and, in each case, the world is as it is represented.

However, we can generate a puzzle by reflecting on the epistemic status of some of the aforementioned truths. For whereas it is clear that we can know that grass is green only by investigation into how things are in the world, it is less clear that investigation of this sort is required to know that there is a prime number greater than twelve or that all vixens are female foxes. Instead, we seem to have the intuition that these are truths

known to us by a little reflection or simply by understanding the meanings of the words involved. The puzzle arises because we now seem to want to maintain both (a) that these are truths by virtue of how things are in the world, and (b) that our knowledge of these truths is able to circumvent investigation into how things actually are in the world. Our knowledge is thought to be *about* the world without in fact being responsive *to* the world. How could such a thing be possible?

One response to the problem is to deny that the a priori truths in question are really true by virtue of how the world happens to be and to hold instead that they do not purport to say anything about the world at all. To the extent that we have a notion of truth apart from considerations about the world, we might say that the sentences that express these are “true in virtue of their meaning.” Considerable disagreement has arisen over the coherence of the view signified by this contentious phrase, but the prototype of the general strategy is commonly supposed to originate in Hume’s distinction between Relations of Ideas – our certainty about which is attributed to their tautological status – and Matters of Fact, the truth of which rests upon the contingencies of the world.¹ This division affords Hume the room to maintain that all substantive knowledge of the world is derived from sense experience, with the a priori truths that would threaten this thesis relegated to the status of mere trivialities which reflect only the meanings we give to certain expressions. In this way, Hume’s motivation for positing the distinction differs in quite important respects from at least some of those responsible for developing the strategy in the twentieth century. Whereas its importance for Hume derives from what we might regard as a foundationalist epistemological programme that seeks to explain the certainty of logical and mathematical truth, its purpose for someone like Carnap is to explain – and indeed to facilitate – the fruitful applications of logic and mathematics to the empirical sciences.² Awodey writes of Carnap:

This distinction, he thought, was not just an artefact of intuition, but had a fundamental practical importance for science. Einstein had said (in a passage often cited by Carnap): “Insofar as the sentences of mathematics refer to reality, they are not certain, and insofar as they are certain, they do not refer to reality...I place such a

¹ Hume 1996, sect. iv

² Cf. Friedman 2007: 4

high value on this conception of geometry because without it, the discovery of the theory of relativity would have been impossible for me.” (quoted in Awodey 2007: 230)

It is fair to say that the strategy of appealing to the idea of “truth in virtue of meaning” has rather fallen out of philosophical fashion. This may be attributable to various factors, though dominant among these are the perception that Quine succeeded in demonstrating the notion to be untenable and, perhaps, Boghossian’s more recent work arguing that there is an alternative conception of analyticity that can do the epistemological work we require which can be separated from one he takes to be laden with the baggage of the failed ambitions of logical positivism. However, it is equally fair to say that there are a number of misconceptions surrounding the idea, including – crucially – what claim the proposal should be understood as being committed to. Boghossian, for example, who labels the view the “metaphysical” conception of analyticity, writes,

Guided by the fear that objective, language-independent, necessary connections would be metaphysically odd, they [the logical positivists] attempted to show that all necessities could be understood in linguistic necessities, in the shadows cast by the conventional decisions concerning the meanings of words. Linguistic meaning, by itself, was supposed to generate necessary truth; a fortiori, linguistic meaning, by itself, was supposed to generate truth. (Boghossian 1997: 336)

But this, he says, is implausible:

Are we to suppose that, prior to our stipulating a meaning for the sentence

Either snow is white or it isn’t

it wasn’t the case that either snow was white or it wasn’t? Isn’t it overwhelmingly obvious that this claim was true *before* such an act of meaning, and that it would have been true even if no-one had thought about it, or chosen it to be expressed by one of our sentences? (Boghossian 1997: 336, italics in original)

Boghossian's comments indicate that he imputes to the proponent of truth in virtue of meaning the view that an analytic truth is one which (i) makes some factual claim about the world, and (ii) makes it the case that the world is that way. Now, I do not know how many, or what proportion, of those who have defended the idea of truth in virtue of meaning Boghossian supposes to have held the position sketched here. But even if some proponents have at times come close to saying something along these lines, we need not regard it as essential to the general strategy that it hold something which so clearly violates our natural intuitions concerning the dependence of truth on the world. For, as we shall see, the claim is not really that there are sentences whose meaning alone is somehow responsible for their truth; rather, the idea is that there are certain propositions which structure our talk about the world and, as such, cannot be regarded as answerable to the world in the same way that an ordinary factual claim answers to the world for its truth. Thus, to the extent that they are not dependent on the world, they do not purport to be about the world.

Another misconception – though one related to that aforementioned – is that the notion of truth in virtue of meaning has in some way suffered from developments made in the philosophy of language over recent decades. Arthur Sullivan, in a recent paper, writes,

One of the reasons why there has been relatively little discussion of TVM [Truth in Virtue of Meaning] in recent decades is that there is such a wide variety of distinct theoretical approaches to meaning. There are orientations that take the basis of meaning to be sense, reference, use, intentions, and truth-conditions.... As a result, it is hard to be any more precise about what a particular claim of TVM comes to..., without alienating a large percentage of the participants in these debates. To proceed in terms of any of sense, reference, use, intentions, or truth-conditions would be controversial; while appeal to all of them would be untidy and vague. (Sullivan 2008: 377)

Taking a similar line, Gillian Russell remarks,

The Quinean camp raised a lot of problems for...analyticity but in the meantime, the "obvious" picture of meaning that supported it started to slip for relatively

independent reasons. In three astonishingly influential pieces of philosophical writing, Putnam..., Kaplan..., and Kripke...[suggested] that the roles attributed to a single thing – the expression’s meaning – ...can be played by distinct things. (Russell 2008: x)³

Sullivan and Russell are surely right to point out that the philosophical conception of meaning has undergone something of a revolution since the idea of truth in virtue of meaning was more prominently advocated, and that we are able to draw on distinctions between content, character, and reference-fixer that would have been alien to earlier proponents of analyticity. In fact, we may even suggest that their comments do not go far enough, since recent work on the semantics/pragmatics distinction has brought to the fore issues about whether the semantic properties of a sentence are sufficient to determine a truth-evaluable content or whether the context of utterance plays an irreducible role in determining what is expressed by a sentence, thereby adding to the issues surrounding our concept of *meaning*.⁴ But it is less clear that these developments should be taken to impact – for better or worse – upon the viability of the idea of truth in virtue of meaning. It may be that we are better able to precisify certain issues by drawing on more recent work, but this is not to say that previous means of articulating the position were significantly deficient or in need of rehabilitation. I hope to show that with a proper understanding of what the view amounts to, we should not regard it as hostage to – or precluded by – any particular conception of meaning.

In section (i) I outline the considerations that can be seen to motivate Carnap’s adoption of the “truth in virtue of meaning” strategy and the role that Carnap envisages these propositions to play within a linguistic framework. In section (ii) I assess the Quinean critique of Carnap’s view, arguing that – contrary to the standard story, as I presented it in the introduction – Quine’s criticisms do not tell decisively against the Carnapian position. In section (iii) I consider a more recent criticism of the so-called “metaphysical” conception of analyticity by Williamson, who argues that the view rests on a conflation of semantic facts with the metasemantic facts. Again, in this case, once we have a sufficient appreciation of what its proponents have advocated, we can see that the resources are

³ Russell is referring, of course, to Putnam (1975), Kaplan (1989) and Kripke (1980).

⁴ See, for example, Recanati (2004), Bach (2005) and Borg (2007).

available to resist Williamson's objection. I conclude in section (iv) by assessing to what extent we should regard the idea of truth in virtue of meaning as in conflict with our natural intuitions concerning the dependence of truth on the world.

(i) Truth in virtue of meaning

There are several routes by which we might account for the development of Carnap's thinking concerning analyticity, which was influenced over a considerable period by the work of Frege and Wittgenstein, as well as developments in geometry, metamathematics and logic. It does not, therefore, lend itself to being presented with any concision, in part because of the sheer number of influences that could be cited as bearing on Carnap's views, but also the complexity of some of these ideas.⁵ Having said that, it is not my aim here to offer anything even approximating a comprehensive overview of these matters. While the historical issues are deeply important for understanding Carnap – and while I hope that what I do say about them here is sufficiently accurate – I shall focus on only a couple of aspects that relate most immediately to illuminating the appeal to truth in virtue of meaning.

The strategy, as we have seen, is to dissolve the puzzle raised in the opening paragraphs of this chapter by holding that there are certain propositions which, contrary to initial appearances, do not purport to say anything about the world and whose truth therefore does not derive from the world being as they represent it to be. Although Carnap was to become perhaps the most sophisticated proponent of truth in virtue of meaning, the general strategy had a contemporaneous forerunner in Wittgenstein. In the *Tractatus*, Wittgenstein expressed a view on which propositions present a "picture" of how things are in the world, their correspondence with which determines their truth-value. A proposition is meaningful insofar as it depicts the world as being some possible way, and we count as understanding its meaning insofar as we know how the world is depicted. As a consequence of this way of thinking about meaning, Wittgenstein argues that there can be no true a priori propositions. For the picture that a proposition presents is just one possible situation in logical space – one way for the world to be – and it is impossible to

⁵ There is, however, laudable scholarship that achieves this; see in particular Coffa (1991) and Awodey (2007).

tell from the picture alone whether it is true or false. An a priori truth would be one whose truth is guaranteed just by virtue what it expresses, “without,” Wittgenstein says, “anything to compare it with.”(TLP 3.05) Such a proposition would therefore not represent the world (since to do so would be to offer something to compare it with), and so would be meaningless.

A proposition can be true or false only in virtue of presenting a picture of the world (4.05). There are, however, certain propositions that do not present any such picture, such as the laws of logic. These, being true in all possible situations, do not present a picture of the world and so not represent any possible situation. Rather, says Wittgenstein, they are the scaffolding that defines the logical space within which we evaluate what is possible (and, therefore, what is meaningful) (3.42, 6.124). They themselves do not say anything meaningful (6.11) precisely because they constitute the boundaries that enable us to say anything meaningful at all.

Notoriously, the initial sense Wittgenstein’s remarks seem to have is undermined when viewed through the very lens they bid us to wear. The claims of the *Tractatus* are not among those that offer a picture against which we can compare the world. But nor are they logical laws of the sort Wittgenstein supposes to define the realm of the factual. According to the very view advocated by the *Tractatus* itself, this is a problem; as Peter Sullivan puts it,

the *Tractatus* holds that ‘there is only *logical* necessity’ (TLP 6.37) – that apart from narrowly logical tautologies the only intelligible claims are contingent. But *that* there is only logical necessity is not something that just happens to be so, and nor is it a tautology. So, if those two are the only kinds of intelligible claims there are, there is no intelligible claim to the effect *that* they are. (Sullivan 2004: 43-4, italics in original)

Sullivan points out that when an analogous problem arose in the 1930s for those who promulgated a verificationist conception of meaning, it was widely taken to be a *reductio* against their view – what would it be to verify the verificationist principle? Wittgenstein, however, seems to sidestep this line of objection by embracing the consequence that the

statements of the *Tractatus* are, strictly speaking, not meaningful. Crucial here is his distinction between that which can be *said* and that which can be *shown*, the former being those propositions that fall within the boundary of the meaningful, the latter being those that lie outside and thereby cannot be expressed. Truths about the role of language cannot be expressed by language (4.121); rather, our use of language shows or reflects those truths. If we were to list all the things that could be said, we would have a complete description of the world. The list would not include the truths about the role of language itself, but these would be shown or reflected in the collection of things that could be said. The strategy of the *Tractatus*, then, as Wittgenstein ultimately admits (6.54), is that we should recognise his remarks as a kind of nonsense, since it fails to be meaningful by its own lights. Yet by grasping what it is that makes it nonsense, one begins to understand the point Wittgenstein wanted to make all along; “one somehow appreciates that there is only logical necessity, and that *therefore* one cannot express that there is; one is *shown* that it is so, and accepts the consequence that one cannot *say* what one is shown.” (Sullivan 2004: 44, italics in original)

The general strategy of the *Tractatus* was a considerable influence on Carnap’s thoughts about the a priori, a fact acknowledged by Carnap in his intellectual autobiography.⁶ Carnap’s mature view, however, grew less directly out of Wittgenstein’s approach to meaning in the *Tractatus*, which he came to reject, and more immediately as a response to concerns about the foundations of mathematics. The problem arose through the challenge faced by the axiomatic method, in which certain propositions are accepted without proof and take the role of axioms or postulates from which all other propositions of a system can be derived. The appeal of employing this method in mathematical inquiry lies in the fact that if one were to establish the truth of the axioms, the truth and mutual consistency of the theorems would be guaranteed. Conversely, if one can derive an inconsistency among the theorems that follow from the axioms, then one has reason to call the axioms into question. It was therefore a source of consternation for mathematical logic that serious contradictions had been uncovered in Frege’s original system and naïve set theory (such as Russell’s infamous class of all classes that are not members of themselves) and that no agreement for how these should be resolved was forthcoming. As a result of this, there

⁶ Carnap 1963a: 25

emerged rival schools of thought with fundamentally different approaches to the foundations of mathematics, with competing views about the direction to be taken in order to respond to the difficulties presented to the axiomatic method.

Against this background, Gödel demonstrated that it is impossible for there to be a consistent formal system containing a certain amount of arithmetic in which every true statement expressible in the system is deducible from its axioms. Even if G is a true statement of such a system that is not formally deducible within it, and we attempt to remedy this by adding G as an additional axiom of the system, the new set of axioms would be guaranteed to yield some further true statement not deducible within the system. The consequence of the result is that for a large class of systems, including elementary arithmetic, it is impossible to prove within the system itself its own logical consistency, unless one adopts principles of reasoning whose proof-theoretic strength leaves them as open to doubt as that of the systems themselves. However, the wider consequence of Gödel's discovery is that, as Nagel and Newman put it,

We are thus compelled to recognise a fundamental limitation in the power of the axiomatic method. Against previous assumptions, the vast continent of arithmetical truth cannot be brought into systematic order by laying down once for all a set of axioms from which *every* true arithmetical statement can be formally derived. (Nagel and Newman 1958: 94, italics in original)

So in the historical context, while there were different approaches to the foundations of mathematics, it would be impossible for any one of these to demonstrate superiority over rival approaches by providing a system in which it could be shown that all and only permissible true statements follow from its axioms, and thus no way of establishing which approach is “the right one.”⁷

It is in Carnap's response to these developments that we can see his considered views about analyticity emerge. Motivated by the desire to make the abstract concerns of philosophy, logic and mathematics less about wrangling over intractable disputes and

⁷ I am extremely grateful to Marcus Giaquinto for improving the accuracy of my remarks about this issue and correcting numerous mistakes on my part.

more about serving the needs of empirical science, Carnap drew a distinction between the framework that structures a system and the useful ends to which that system may be put. Just as Wittgenstein in the *Tractatus* had distinguished between the meaningful sentences that present a picture of the world and the strictly meaningless claims which nevertheless define the realm of the factual, so Carnap in similar fashion separated the empirically significant statements of a system from those which serve to structure – and give meaning to – these empirically significant statements.

This way of drawing the distinction was crucial for Carnap. Whereas for Wittgenstein the laws of logic constituted the very boundary of thought itself (*TLP* 4.114), for Carnap our adoption of a set of logical rules is a pragmatic affair, determined by what best serves the empirical ends it is used to pursue. This was dubbed by Carnap the “Principle of Tolerance”, though as commentators are keen to point out, it is in fact less a principle and more an attitude,⁸ one born of the Wittgensteinian idea that logical truths do not describe the world but instead structure the language used to describe the world. Since they are not really about the world, no question about their truth or objective correctness arises, and so we are therefore free to adopt the principles that best suit our needs. Articulating this stance in *The Logical Syntax of Language*, Carnap writes,

let any postulates...and rules of inference be chosen at will; this choice gives the logical signs occurring in them the meaning they have. With this attitude, the conflict between different viewpoints on the problems of foundations of mathematics disappears. The mathematical part of the language can be set up as one or the other viewpoint prefers. There is no question of “justification,” but rather only the question of the syntactic consequences to which one or another choice leads, including the question of consistency. (quoted in Ricketts 2007: 203-4)

We might be tempted to question the accuracy of Carnap’s claim to have obviated the potential for disagreement about the foundational issues of a system, observing that room for disagreement is removed only if all frameworks are judged to be of equal merit, which is not the view Carnap himself promoted. In both earlier and later works, he cites such

⁸ See Ricketts 2007: 218 and Creath 2007: 322

factors as the purpose of the system, its efficiency, fruitfulness, and simplicity of use as considerations material to the adoption of a particular framework.⁹ To the extent that we may disagree about the fruitfulness or simplicity of one framework over another, there remains scope for disagreement about which we should adopt for any given purpose. But this would not, in the historical context, necessarily be of concern to the overarching Carnapian outlook, since it nevertheless succeeds in replacing what he would regard as fruitless “metaphysical” argument about which is the objectively correct way of doing things with a more tangible, scientifically grounded criteria of practical utility. For Carnap, these issues of simplicity and fitness for purpose move the disagreement out of the grip of purely abstract theorizing into the domain of the empirically verifiable.¹⁰

Carnap’s view, in certain significant ways, can be seen as a precursor to the now familiar idea made famous by Wittgenstein’s later, post-*Tractatus* writing of a language system being like the rules of a game. The analogy is cashed out by Ricketts in the following way:

The choice of a formal language as the language for science...resembles the choice of rules for a game we intend to play. The rules define what is permitted in the course of the game. But no such question of legitimacy applies to the choice of these rules themselves. Rather, we construct the rules for the game by weighing various considerations in order to frame rules that define, say, an enjoyable, engaging, competitive activity whose course depends on a mixture of physical skill, quick wits, strategy, and luck. We may, in the light of our experience playing the game, decide to modify the rules, perhaps in far-reaching ways, and so to play a different game. Similarly, the option to change languages is always open as regards the choice of a language for science. (Ricketts 2007: 207)

The aim of the language is to be able to say things about the domain with which the language is concerned, whether that be an area of mathematical enquiry or some area of empirical investigation. Making a meaningful statement concerning the objects of the domain constitutes a move within the game, as laid down by the rules that define the activity. These rules are neither determined arbitrarily nor is there any sense in which

⁹ Carnap 1956b: 208; Ricketts 2007: 206-7

¹⁰ Cf. Friedman 2007: 9

they could be objectively correct. They are, rather, designed to best facilitate the overall aims of the practice, viz. to say things about the relevant domain of objects. In the course of playing the game, participants might decide that the current rules do not optimally facilitate certain aspects of its aim. As a result, they may opt to alter the rules and thereby change, to some degree, the activity which these define.

There are, then, for any given system, two different types of statement that might be made: those which purport to say something about the objects of the domain, and those which articulate the rules governing the system and by virtue of which statements of the previous kind can be made. Among the latter Carnap counts, among others, the rules of designation for determining the reference of individual constants and predicates, the rules of formation and transformation that determine the ways in which expressions can be meaningfully combined with one another, the statements that implicitly define the basic logical constants, and the basic rules of inference. With these rules in place, we can construct what Carnap, in *Meaning and Necessity*, calls a state-description for the domain of a language. A state-description is a class of sentences which contains for every atomic sentence in the language either the sentence or its negation, but not both, and thus provides “a complete description of a possible state of the universe of individuals with respect to all properties and relations expressed by the properties of the system.” (Carnap 1956a: 9) In this way, Carnap’s talk of state-descriptions is akin to talk of possible worlds, though the possibilities are determined not by any metaphysical facts about the world (as some talk of possible worlds suggests) but rather by the rules laid down for the particular system.

State-descriptions, like possible worlds, are mutually exclusive of one another: if one represents the true or actual state of affairs, then no others do. A sentence holds in a given state-description if it would be true were that state-description to be the true or actual one. The class of state-descriptions in which a sentence *S* holds is what Carnap calls the range of *S*. The range of a sentence is determined by the rules of the system. A sentence whose range is less than the total number of state-descriptions is one which is capable of being true or false; investigation is required into how things happen to be in order to ascertain its truth-value. These are the synthetic truths of the system. A sentence whose

range is equal to the total number of state-descriptions is analytic. The truth of such a sentence can be established on the basis of the semantical rules of the system alone, without any reference being required to extra-linguistic fact. For example, we know that whatever the range of some atomic sentence *S*, the state-descriptions not containing *S* contain not-*S*; consequently, we can establish that the disjunction *S* or not-*S* holds for all state-descriptions and that, whichever state-description happens to represent the actual state of affairs, *S* or not-*S* must be a truth of the system.

We can now begin to see a strategy by which to dissolve the puzzle set up at the beginning of this chapter. If we can draw a distinction between the substantive statements of a language, which have a factual content about the objects in the relevant domain, and the statements which constitute a framework that structures the substantive statements, we thereby admit a class of truths which do not lend themselves to assessment for truth or falsity by how things are with the objects of the domain. These are the statements that function as rules of the system. As such, they are not to be regarded as having the same status as the substantive truths of the language; they do not purport to make factual claims about the domain nor, thus, should their truth be supposed to derive from how things are within that domain.

The error in the prior characterisation of the strategy by Boghossian should be clear. Meaning does not generate truth, at least not in the sense implicit in Boghossian's remarks, where an analytic truth is one that both (i) makes some factual claim about the world and (ii) makes it the case that the world is that way. This claim – Boghossian is surely right to insist – is absurd. However, as we have seen, proponents of the strategy deny (i) (and, to my knowledge, have said little by way of countenancing (ii)). They have instead argued that the constitutive statements of a language are held true by the users of a language because they have the status of structuring the very language in which they participate. The epistemological upshot is that anyone who understands the language has grounds to endorse such a statement without having to check whether the world complies with its truth; a grasp of the rules by virtue of which the expressions of the language have their meaning is sufficient to enable one to see that its truth is guaranteed.

(ii) Carnap and Quine

The dispute between Carnap and Quine is probably the most famous chapter in the history of the notion of analyticity. Yet it remains, as we shall see, a particularly difficult debate on which to adjudicate. Although many have considered Quine's arguments to tell decisively against the Carnapian account of analyticity, it is less common to find the disagreement reconstructed in a way that does justice to its principal players. The difference between their respective views is so slight that Carnap went so far as to imply that Quine's position was tantamount to his own, yet so fundamental that perceptive commentators have suggested that "there appears to be no way even to judge what kind of dispute it is without thereby taking a side in it." (George 2000: 19)

Quine's basic outlook is shaped by a thoroughgoing scientific empiricism, according to which our philosophical claims about the world should be continuous with what we are told about how the world is by the natural sciences. The optimal grounds for holding any statement to be true are that it either falls directly within the limits of scientific observation or otherwise belongs to the collection of statements which comprise our best scientific account of the world. Thus, for Quine, how we find the world to be has an essential role for all statements we are prepared to countenance as true; there can be no statements held true independently of how the world is. Even the laws of logic we employ derive their truth from the general features of the world, as he writes in 'Carnap and Logical Truth':

Another point...was that true sentences generally depend for their truth on the traits of their language in addition to the traits of their subject matter; and that logical truths then fit neatly in as the limiting case where the dependence on traits of the subject matter is nil. Consider, however, the logical truth 'Everything is self-identical', or $(x) (x = x)$. We *can* say that it depends for its truth on traits of the language (specifically on the usage of '='), and not on traits of its subject matter; but we can also say, alternatively, that it depends on an obvious trait, viz., self-identity, of its subject matter, viz., everything. (Quine 1966c: 113, italics in original)

The sentiment expressed in the above passage is often cited approvingly by those who take Quine to have decisively refuted the idea that there could be such a thing as truth in

virtue of meaning:¹¹ all statements held as truths are so because they are regarded as answering to the world in the appropriate way. It is not clear, however, that the mere assertion of this stance does much to engage at a critical level with the Carnapian view to which it is supposed to stand in opposition. For it is open to Carnap to reiterate that insofar as a statement is *about* the world, it must be true by virtue of how things are *in* the world, yet to maintain that this very practice of speaking about the world presupposes certain rules which are not themselves about the world and therefore not answerable to the world. Moreover, from the Carnapian perspective, the idea that we can say – as Quine suggests – that the truth that everything is self-identical derives from our concept of identity, and not on specific features of the world, indicates that this must be among those statements whose truth is guaranteed by the rules of the language. In doing little to cut off this response, the quotation from Quine may be seen as an articulation of his general viewpoint rather than as an argument in its favour.

An appreciation of Quine's outlook is nevertheless important, since the principles he employs in arguing against analyticity emanate directly from this. His view that empirical statements should fall within the limits of scientific observation underpin a methodological behaviourism, according to which the admissible statements about a subject are those available from the perspective of a third-party investigator. In attempting to record the facts about what a person means, believes or intends, we are to eschew the idea that there is knowledge that would aid or illuminate this effort that is available only to the subjects under observation, since this would seem to have the consequence of putting these facts beyond the reach of independent scientific enquiry.¹² Allied to this is an epistemological holism, according to which the totality of statements about the world to which we subscribe is to be assessed as a whole for its overall fit with the body of evidence at our disposal and their coherence with one another. This totality forms a "web of belief" in which the interconnections between the various strands mean that revising our acceptance of one statement involves re-evaluating, and perhaps modifying, others elsewhere in the system. For Quine, this is a natural accompaniment to

¹¹ See, for example, Boghossian 1997: 335, Margolis and Laurence 2001: 294, and Russell 2008: 29. Although Russell does hold that there is a viable notion of truth in virtue of meaning, hers is not that found in Carnap, which she does suppose to be subject to the Quinean criticism.

¹² Cf. Alston 1986: 67

his scientific empiricism, since it is always possible that observation will yield evidence that forces us to revise our account of the world. These revisions are always constrained by the aim to find the best fit between the various statements we endorse. As a result, there is no statement that can be immune from revision if revising it produces the best fit between the other statements within the web.

We can see both the principles of methodological behaviourism and epistemological holism at work in the two main lines of argument Quine offers against the tenability of an analytic-synthetic distinction in his famous paper ‘Two Dogmas of Empiricism’.¹³ These two arguments correspond to two different types of response that might be offered by way of providing a criterion for distinguishing between analytic and synthetic truths. The first is the attempt to explicate the notion of analyticity in other terms. For it is quite natural to suppose that certain paradigmatic examples of analyticity, such as conceptual truths like ‘All bachelors are unmarried men’ or ‘All vixens are female foxes’, can be explained by appeal to a sameness of meaning between particular words and phrases, or perhaps their having the status of definitions or necessary truths. But this is at odds with Quine’s behaviourism, since all admissible claims about a subject are those available to the observation of a third-party spectator. When we apply this stance to the study of meaning facts – what a subject means by the use of particular words – we are thus restricted to the correlations we can observe between the subject’s usage and the objects or circumstances that appear to cause or prompt that usage – with, of course, a holistic approach to working out the network of meaning had by the totality of a subject’s words. This approach precludes, *ipso facto*, the kinds of intensional properties that would be needed in order to be able to appeal to sameness of meaning. For we do not suppose that mere extensional agreement suffices for sameness of meaning; the phrases ‘creature with a heart’ and ‘creature with a kidney’ may have the same extension, though we do not suppose they have the same meaning. The problem for analyticity is that extensional agreement is all observation can deliver: we may establish that ‘bachelor’ and ‘unmarried man’, ‘vixen’ and

¹³ Precisely how we ought to construe the argument of ‘Two Dogmas’ is itself a matter of some contention, as illustrated by the survey of differing interpretations presented in Creath (2007, esp. 327). I have here put aside the finer points of Quinean scholarship for present purposes and offer my own understanding of the text, which coheres with the interpretations offered by Creath (1997) and George (2000), as well as much of the additional writing between Quine and Carnap, as collected in Creath (1990).

'female fox', have the same extension, but this does not show that they have the same meaning.

Synonymy, then, is unable to ground the notion of analyticity. We might invoke another concept instead. Quine raises several alternatives, including *definition*, *interchangeability*, *necessity*, and *semantical rule*, arguing that they are all part of the same mutually supporting circle of interrelated concepts. The problems that affect one infect them all. For while Carnap may insist that we know which are the semantical rules of a language (they are those statements we adhere to in forming the meaningful sentences!), this would require that some specification be offered among the rules to the effect that they have this status. But how, asks Quine, can *this* be among the rules if we do not already understand what status *this* is supposed to signify?

The rules tell us that such and such statements, and only those are the analytic statements of L₀. Now here the difficulty is simply that the rules contain the word 'analytic', which we do not understand! We understand what expressions the rules attribute analyticity to, but we do not understand what the rules attribute to those expressions. (Quine 1953: 33)

Each of the expressions within the circle is explicable only in terms of another. Appeal to any of them will not enable us to draw a distinction between the analytic and the synthetic.

The second line of argument concerns whether there could be any behavioural evidence adduced in support of the distinction. Given Quine's broader outlook, this would clearly constitute an acceptable form of evidence in favour of a division between the analytic and synthetic statements of a language, were any to be offered. Yet it is not immediately obvious what sort of behavioural evidence would support the distinction. Quine's remarks suggest that it would be only if there were evidence that certain statements of a subject's language had to be held true "come what may", in spite of any experiences perceived as presenting evidence to the contrary. But, of course, there are for Quine no statements that could fall into this category in view of the holistic approach that governs our acceptance

of scientific hypotheses. On this approach, he argues, we do not consider evidence on an individual statement-by-statement basis; assessment applies to the collection of statements as a whole, which we may amend any way we require in order to get the best fit. There are, then, no statements that could ever be held come what may – at least, not in a way that would distinguish them from any other statements of the system. It is open for us to maintain any statement of the system, just so long as we are prepared to make the appropriate adjustments in the system; he writes,

Any statement can be held true come what may, if we make drastic enough adjustments elsewhere in the system. Even a statement very close to the periphery can be held true in the face of recalcitrant experience by pleading hallucination or by amending certain statements of the kind called logical laws. Conversely, by the same token, no statement is immune to revision. Revision even of the logical law of the excluded middle has been proposed as a means of simplifying quantum mechanics; and what difference is there in principle between such a shift and the shift whereby Kepler superseded Ptolemy, or Einstein Newton, or Darwin Aristotle? (Quine 1953: 43)

So, no statements are in principle of a kind that requires them to be held true come what may. To the extent that we *can* choose to hold a given statement true, it does not differ in that respect from any other statement, which could equally be held true. There is, therefore, no behavioural evidence that could ground a distinction between the analytic and synthetic statements of a system.

On the face of it, Carnap has fairly straightforward responses to offer the Quinean objections. In reply to the first, he can accept the interdefinability of the relevant expressions, without having to accept that the circularity is vicious. Once we have reason to posit one member of the familial set – such as that of the semantical rule for the kinds of consideration seen the previous section – we have a way into the circle that begins to make sense of them all. Viewed this way, there is no problem with explicating one expression in terms of another. This point is put to Quine by Carnap in his response to ‘Carnap and Logical Truth’, in which he argues that Quine cannot really be objecting to the circularity here since he has no similar reservation about the concept of *truth*, even

though we could only explicate it by appeal to a similar family of interdefinable concepts.¹⁴ Carnap says that Quine's objection boils down to the perceived lack of a sufficiently clear explicandum for the concept *analyticity* (in which respect it differs, for Quine, from *truth*), which is the reason that leads Quine to demand an empirical criterion for the notion – which brings us, of course, to Quine's second objection.

Although Carnap holds that there *is* empirical evidence that can be cited in support of the claim that a particular sentence is analytic in a given language (and outlines the sort of test that might confirm or disconfirm the hypothesis),¹⁵ he does *not* believe that there is empirical evidence to be adduced in support of drawing the distinction between analytic and synthetic truths. Responding to Quine, he writes,

The main point of his criticism seems...to be that the doctrine is “empty” and “without experimental meaning”. With this remark I would certainly agree, and I am surprised that Quine deems it necessary to support this view by detailed arguments. In line with Wittgenstein's basic conception, we agreed in Vienna that one of the main tasks of philosophy is clarification and explication. Usually, a philosophical insight does not say anything about the world, but is merely a clearer recognition of meanings or of meaning relations. If an insight of this kind is expressed by a sentence, then this sentence is, although meaningful (as we would maintain in contrast to Wittgenstein's view), not factual but rather analytic. Thus I would interpret, e.g., the principle of verifiability (or of confirmability), or the empiricist principle that there is no synthetic a priori, as consisting of proposals for certain explications (often not stated explicitly) and of certain assertions which, on the basis of these explications, are analytic. (Carnap 1963b: 917)

Carnap's point is to remind us that the original purpose of the logical positivists' appeal to the analytic-synthetic distinction was, in part, to eliminate what they saw as the fruitless questions of metaphysics by dividing the statements of a language into two kinds: those which are empirical claims about the world (and thus subject to empirical confirmation or disconfirmation) and those which reflect the rules governing the meaningful things that

¹⁴ Carnap 1963b: 919. He does not actually state the other concepts which he supposes to comprise this particular circle, but presumably he would cite such concepts as *meaning*, *facts*, etc.

¹⁵ Carnap 1963b: 920

can be said, which are not themselves about the world and so not answerable to it by empirical investigation. Philosophy, when it falls into the realm of metaphysics, becomes meaningless, since it falls into neither of these camps; done properly, on the other hand, it serves to clarify the conceptual issues that structure our empirical claims. The claim that there is a division to be drawn between analytic and synthetic statements in a language aims to clarify an aspect of our linguistic practice and is, hence, a philosophical one. It is therefore no surprise that it should have no empirical criterion by which to confirm it – on this view, no philosophical statements do! That we can draw a distinction between analytic and synthetic truths is itself, Carnap seems to want to say, an analytic truth.

We need to be careful about how this point is made. Although I have put Carnap's point this way in order to convey his viewpoint, he does not – so far as I know – ever state it quite so explicitly. And there is good reason for this. It is important to remember that, for Carnap, analyticity is always relative to a language. The point of viewing analytic statements as factually empty was that we should be able to adopt a different set as it proved convenient for a language of science, though when we do so we are thereby adopting a different language. The problem with saying 'That there are analytic truths is analytic' is that it can only be analytic relative to some language, which would mean we could always change our language so that it is no longer analytic, whereas Carnap's view is better understood as the idea that it is a language-transcendent, necessary truth that there *must* be analytic truths. But, just as with the bind that prevailed on Wittgenstein's metaphilosophical discourse, this looks like something which – by its own lights – cannot actually be said but must rather be appreciated in silence.

It is tempting for those instinctively sympathetic to the notion of analyticity to see all this as a fairly straightforward vindication of the Carnapian position against Quine's objections. After all, in pointing out the interdefinability of 'analyticity' and related expressions and the absence of an empirical criterion for drawing the analytic-synthetic distinction, Quine appears to have offered nothing which Carnap does not already accept nor any reason why these considerations should tell against the notion of analyticity. However, there is good reason to think that this cannot be a conclusive analysis of the debate. As George points out, we need only reflect upon Quine's basic outlook to recognise his grounds for

being dissatisfied by Carnap's reply.¹⁶ Quine regards it as a condition upon a properly scientific approach to philosophy that it be viewed as continuous with the natural sciences and, moreover, that experience should have an essential role for all statements we are prepared to accept as true. Carnap, by contrast, admits a role for philosophy which he supposes to contribute to the scientific enterprise, by clarifying the conceptual issues which structure our empirical claims, but which must be distinguished from the doing of the science itself insofar as it is not subject to empirical investigation. Their disagreement over the tenability of the analytic-synthetic distinction can be traced to these very different starting points. Carnap, we have seen, takes the claim that there are analytic truths to be a philosophical one and, as such, in no need (as well as incapable) of empirical support. For Quine, the admission that the claim is without empirical support is tantamount to saying that there is no reason to suppose it true. This, George argues, is why we find Carnap agreeing with Quine on the points that Quine intends as objections against analyticity:

Quine's arguments...make a point about evidence, namely that the [analytic-synthetic] distinction is an empirically empty one. When conjoined with Quine's other beliefs, this leads to a rejection of the distinction. But when conjoined with Carnap's, the point functions instead as a ratification of his conception of the distinction. (George 2000: 9)

Quine, given his own starting point, sees no reason for Carnap to remain wedded to the notion of analyticity. There is not, he feels, anything to be gained by subscribing to the analytic-synthetic distinction that cannot be accommodated within his own system; so someone like Carnap, who professes to desire a genuinely scientific approach to philosophy, should abandon what Quine supposes can only be a dogma. For his own part, Carnap suggests that there is more agreement between the two than Quine seems able to appreciate, and that if Quine were to reflect upon his own metaphilosophical commitments, he would recognise the common ground which they share. Making the suggestion somewhat elusively in a response to Quine, he says,

¹⁶ George 2000: 7-8

I believe that that the distinction between analytic and synthetic statements, expressed in whatever terms, is practically indispensable for methodological and philosophical discussions. This is also indicated by the fact that this distinction is made by a large majority of philosophers, including some of those who do not explicitly acknowledge the distinction in these terms or even reject it. As an example, let me refer to a philosopher whose work I esteem very highly, although I cannot agree in all points with his views. This philosopher once undertook to destroy a certain doctrine, propounded by some other philosophers. He did not mean to assert that the doctrine was false; presumably he regarded it as true. But his criticism concerned its particular kind of truth, namely that the truth of the doctrine was of the analytic kind. To be sure, he did not use the word “analytic”, which he did not seem to like very much. Instead, he used other expressions which, nonetheless, clearly seem to have essentially the same meaning as “analytic”. What he showed was that various attempts to assign an experimental, empirical meaning to this doctrine remained without success. Finally he came to the conclusion that the doctrine, even though not false, is “empty” and “without experimental significance”. (Carnap 1963b: 922)

The idea here, I take it, is driven by Carnap’s conviction that every practice comes with a framework, with its own set of rules or principles which cannot be shown within the system to be objectively correct. These rules or principles constrain the statements that can be admitted as true within the practice. Presumably for Quine these would include the idea that philosophy should be continuous with natural science, or that all statements answer to our experience of the world for their truth. Having subscribed to these principles, Quine is unable to admit of any statements that could be held as true without empirical support. But, Carnap implies, Quine should not really regard himself as committed to that stance if were to reflect upon the status that his own principles have within his philosophical system. Quine, for his own part, can deny that his principles are lacking in empirical content in the way Carnap supposes of framework propositions, maintaining instead that they are as much subject to empirical consideration as any other statements we hold true, to be assessed according to the experiential criterion of the fruitfulness with which the resulting methodology yields knowledge about the world.

How, then, are we to adjudicate on such a debate? The problem is that there seems to be no way to decide on the criterion by which to judge the disagreement that does not presuppose a particular stance within the debate itself.¹⁷ If we think that it is a philosophical question that can be answered without any appeal to our experience, and without any consideration for the needs of scientific practice, we have already adopted a position broadly in sympathy with the Carnapian perspective. If, on the other hand, we suppose that considerations about our experience are necessary here, then we would seem to be disposed to favour the Quinean view.

So, do we have a genuine philosophical impasse? Perhaps, though we would need a more thorough examination of the relevant issues than we have even come close to here in order to establish such a conclusion. Certainly, however, it is difficult to see which resources at the disposal of one theory could be deployed as telling against the other without merely reasserting the differences between them. Yet even the view that there is ultimately little to be said to favour one over the other contains a hint of the Carnapian position, George suggests, given the considerations we have seen to motivate his stance.¹⁸

Should we regard this as a problem? I do not suppose that many philosophers are inclined to be content with an unresolved stand-off between two accounts that in certain respects are so different from one another. Where there is a truth of matter to be had as to which, if either, is correct, it is natural to desire that we should know what it is. To this extent, we might take the prospect of an impasse to be a problem. Within the much narrower context of the present considerations, however, it may legitimately be regarded as less of a problem. The dialectic, as it has been presented here, is one concerned with the possibility that there could be truths evidentially disconnected from experience. We started from the intuition that there does appear to be such a class of truths and have observed one strategy by which this might be explained. Our interest in the dispute between Carnap and Quine stems from the fact that has been widely viewed as enormously significant, not simply in being a long-standing disagreement between two figureheads of twentieth-century philosophy, but also because many subsequently adjudged Quine to have struck a decisive

¹⁷ The point is made by George 2000: 17-19.

¹⁸ George 2000: 18-19

blow against Carnap's idea of truth in virtue of meaning and any other analytic explanation of the a priori that might rest upon any similar notion. But this, we can now see, is not a conclusion we need be compelled to accept. It may be that there are aspects of the Quinean outlook that have made it the more attractive for some of a certain philosophical temperament to endorse, but on what we have seen, the arguments themselves are not decisive.

(iii) Williamson and Benacerraf on semantic uniformity

More recently, Williamson has raised a different line of objection against the attempt to explicate analyticity in terms of truth in virtue of meaning. His strategy, to this end, is to challenge the move from the fact that we can stipulate rules for our expressions, so as to guarantee the truth of certain sentences, to the idea that those sentences are thereby not about the world or answerable to it for their truth. We can make it the case that particular expressions *mean* what they do, but we cannot make it the case that the world *is* any particular way. To suppose otherwise, he suggests, is to confuse the semantic facts of a language with the metasemantic facts of the language.

The semantic facts are those concerning what the expressions of a language mean, which objects, properties and relations they signify. The metasemantic facts are those concerned with how the expressions of a language come to have the meaning they do; these are non-semantic facts upon which the semantic facts supervene. For example, that the name 'Dubya' denotes the man who became the forty-third President of the United States is a semantic fact; that it came to do so as a convenient shorthand by which to distinguish references to him from those to his all too similarly named father is a metasemantic fact. Similarly, that 'Hesperus' denotes the same object as 'Phosphorus' is a semantic fact; that it came to do so because the ancients did not recognise the identity of the heavenly body most prominent in the evening with that most prominent in the morning is a metasemantic fact.

The significance of the distinction lies in the fact that insofar as we are interested in the truth of a statement, it is only the semantic facts which are relevant. The metasemantic facts determine how the constituent expressions came to have the meaning they do, but

once the meanings of the expressions are assigned, a statement is meaningful by virtue of combining them in such a way as to represent the world as being some way or other. And once it does this, we can assess it for truth or falsity according to whether or not the world is as it represented to be.

Williamson illustrates the point by considering the introduction of the expression ‘zzz’ (pronounced ‘buzz’) to mean a short sleep. In this instance, that ‘zzz’ means a short sleep is a semantic fact; that it is introduced by an act of stipulation that it should do so is a metasemantic fact. Williamson writes,

My saying ‘A zzz is a short sleep’ did not make a zzz be a short sleep, because that would be to make a short sleep be a short sleep, and my saying ‘A zzz is a short sleep’ certainly did not make a short sleep be a short sleep. In particular, since there were many short sleeps before I was born, there were many zzzes before I was born, independently of my later actions. At best, my saying ‘A zzz is a short sleep’ made ‘zzz’ mean a short sleep, and therefore ‘A zzz is a short sleep’ mean that a short sleep is a short sleep. This is simply the standard semantic contribution of meaning to truth, just as for synthetic truths. (Williamson 2007: 82-3)

So, according to Williamson, stipulating rules for a language does not deliver truth in virtue of meaning of the kind necessary to ground analyticity. We can make pronouncement about what particular expressions are to mean, for example, by deciding at some point that ‘vixen’ is to mean ‘female fox’. And this might have the consequence that certain sentences are guaranteed to express truths, as ‘All vixens are female foxes’ does since there is no possible world in which the sentence could not be true, given the meanings assigned to the expressions in question. But these expressions are meaningful by virtue of the fact that they denote some way for the world to be, and it is the world’s being this way that makes the sentence true, not our stipulations about the meaning of the expressions.

Now, it should be clear that the conception of truth in virtue of meaning against which Williamson takes himself to be arguing is one deeply influenced by the characterisation we saw earlier to be offered by Boghossian. Williamson supposes that the advocate of

truth in virtue of meaning intends to be understood as claiming that our meaning what we do by our words somehow makes it the case that the world is one way or another. This misconception of the strategy leads Williamson to present an argument against it that is simply question-begging. For Williamson argues in effect that ‘A zzz is a short sleep’ is not true in virtue of meaning on the grounds that it is semantically equivalent to ‘A short sleep is a short sleep’; and since it is quite clear, he supposes, that the latter answers to the world for its truth, so must the former also. But the latter is an instance of a logical law of precisely the kind which, advocates of truth in virtue of meaning argue, does not answer to how things are in the world.

Why, then, we might wonder, should any notice be taken of an argument that so badly misconstrues its target view? Because, I want to suggest, it is an argument that exhibits symptoms of a more general concern that one might raise against the idea of truth in virtue of meaning. The strategy we saw in both Wittgenstein and Carnap was to divide the statements of a language into two kinds: those which are meaningful, empirically investigable statements that answer to the world and those which provide the structure within which the aforementioned have meaning and so do not themselves answer to the world. In doing this, they challenged what Coffa describes as the “semantic uniformity” of statements. Expounding upon this claim, Coffa says,

Those before them who had attempted to explain the character and ground of a priori knowledge had considered the fine structure of judgment as well as any distinctions between the semantic tasks performed by different kinds of statements, to be irrelevant. For example, for Kant a priori and posteriori judgments were species of a semantically uniform genus, *judgment*; therefore, everything that “Logic” told us about judgment would, ipso facto, be applicable to the a priori. According to this picture of things, in order to understand the difference between the a priori and everything else we would have to go beyond the study of judgment to a transcendental theory of mental activity; the strategy was to look at how we can *know* such claims rather than at what it is that they *say*.

(...)

Wittgenstein and Carnap argued that the main difference between a priori statements and the rest lies at the level of what they say. A priori statements, they asserted, do

not differ from others merely in the strength with which the mind or other things entitle us or force us to believe what they say. It lies rather in the extraordinary role these statements play in the process of making it possible to say anything at all. (Coffa 1991: 259-60, italics in original)

The claim of semantic uniformity is that we treat all statements of a language as on a par in their semantic assessment; there are not, on the one hand, statements we can say are about the objects of the domain and, on the other, statements to which we can assign the status of rules. And something of this claim seems to underpin Williamson's complaint against truth in virtue of meaning. Once the meanings of the individual expressions have been determined, we assess the statements in which they feature equally, that is, according to the criterion of whether the world is as it represented to be.

I take it that the claim, when we consider it, has an obvious appeal. In the opening paragraphs of this chapter, various examples of declarative sentences were listed, some putatively analytic and some quite clearly not. I claimed then that we have a common intuition that they are all equally about the world. We do not "hear" a difference between them that puts some into a group that are simply about the world and some into another group that are rules. So, although neither Wittgenstein nor Carnap are committed to the view that the division they propose should be in any sense obvious or phenomenologically manifest to individual speakers, it is nevertheless a cost to their rejection of semantic uniformity that it should posit such a division where none is apparent – one which must be weighed against the explanatory benefits the strategy offers.

The worry echoes Benacerraf's paper 'Mathematical Truth', in which he argues that in the case of mathematical truth (though with application to the a priori more broadly) there seems to be an unfortunate trade-off between having a reasonable epistemology and a homogenous semantic theory where the semantics of mathematical statements parallel the semantics for the non-mathematical statements of a language. Benacerraf suggests that "almost all accounts of the concept of mathematical truth can be identified with serving one or another of these masters *at the expense of the other.*" (Benacerraf 1973: 661, italics in original) The dilemma he poses is one on which we can have a uniform semantic

analysis of both mathematical and non-mathematical statements, perhaps by treating the objects of mathematical statements as being “out there” just like the objects of ordinary discourse, but consequently struggle to account for how we have knowledge of such objects; or we can have a less extravagant epistemology for mathematical truths, perhaps grounded in human practice or conventions, and settle for a disparate semantic analysis of these statements, since they no longer depend on the satisfaction of a condition upon the world for their truth.

For the defender of truth in virtue of meaning, however, Benacerraf's dilemma is a false one: there is no genuine trade-off to be had because there is no genuine option of preserving semantic uniformity. There is no suggestion in Wittgenstein and Carnap that giving up uniformity is just the unfortunate price to pay for having a palatable epistemological account of the a priori. Rather, we saw it argued that the very notion of a semantical system that allows us to say things about the world presupposes a set of rules by which the expressions have their meaning. The statements that articulate these rules constitute a framework for the system within which meaningful activity can occur. Thus, semantic uniformity is neither accompaniment nor accomplice to one epistemological account or other. It is instead a claim that Wittgenstein and Carnap suppose to conflict with the possibility of there being meaning at all. On this view, then, to forego an analytic explanation of the a priori in favour of a rationalist one for the sake of preserving a thesis of semantic uniformity is the philosophical equivalent of being led further into a desert in pursuit of a mirage.

Yet this claim, that it is simply not possible to provide a unified semantic analysis of all statements in a language, certainly appears open to question. After all, one might argue, surely that is precisely what we offer when we construct semantic theories on which the meaning of sentences in a language are identified with their truth-conditions? We can deploy a disquotational schema of the form ‘S is true iff p ’ which delivers for any sentence on the left-hand side of the schema a condition upon the world that is to be met in order for the sentence to be true. This appears not to discriminate between those sentences Carnap classifies as factual and those he supposes to be analytic. We can plug in the sentence ‘Grass is green’ and the disquotational schema tells us that it is true iff grass is

green; but equally, we can input a sentence like ‘All vixens are female foxes’ or ‘Everything is self-identical’ and be given a condition upon the world to be fulfilled, namely, that all vixens are female foxes and everything is self-identical respectively. It thus appears that a unified semantic treatment is available after all.

The defender of truth in virtue of meaning may point out that both the schema and the appeal to truth-conditions presuppose that the sentences in question impose a condition upon the world that is to be satisfied in order for it to be true. In assuming that the putatively analytic truths answer to the world in this way, this attempt to preserve semantic uniformity merely reasserts what the defender of truth in virtue of meaning denies. But what the advocate of semantic uniformity achieves is to highlight again that we do not hear anything intuitively problematic with saying something like “All vixens are female foxes’ is true iff all vixens are female foxes’, just as we might apply the disquotational schema to unarguably empirical claims of the language. We have little difficulty with the thought that it is a truth which is about the world and as having a truth-condition, albeit one that is necessarily satisfied given the meaning of the sentence. We do not hear the contingent, empirical statements as being about the world in a way that the necessary, non-empirical statements are not, with the former as ordinary factual statements and the latter as rules which do not answer to the world. And it is this intuition that the advocate of semantic uniformity can push in opposition to the notion of truth in virtue of meaning.

The defender of truth in virtue of meaning, we have already noted, need not deny that there is an appearance of semantic uniformity; he rather denies that this appearance of uniformity is borne out by the role played by different statements in a language. It seems that this move requires some sort of explanation for why there should be an appearance of semantic uniformity when there really is none. I want to suggest that the defender of truth in virtue of meaning can offer just such an explanation – one which mirrors the treatment of necessity that emerges from Wittgenstein’s and Carnap’s views.

For Wittgenstein and Carnap, there are no statements that express what we might think of as objectively necessary facts about the world, that is, no statements which are about

the world and which express a fact about it that could be true in all conceivable circumstances independently of what we mean by the constituent expressions. Meaningful, factual statements are confined to saying things about the world that could conceivably be true or false. Yet there remain certain statements that appear to express truths which hold 'by necessity', i.e. the analytic statements. Wittgenstein and Carnap, of course, deny that these have any objective truth; they are merely the boundaries which shape our thought about the world. At the same time, for anyone whose thought about the world is shaped by these statements, they will naturally be regarded as necessary since they define the possibilities we are able to conceive of the world as having. Coffa expresses the position nicely, writing,

what people call necessary is indeed necessary when one looks at it, as it were, from the inside; but in order to understand it fully, one must also look at it from the outside – and what one sees *then* is that it is conventional. Moreover, one also sees that these conventions were really definitions in disguise; syntactic appearances notwithstanding, these sentences have more to do with identifications of meanings than with the making of claims. The statements that constitute what Carnap would...call a language framework are, within that framework, necessary; that is, to the extent that we are using that system, we cannot deny them. But we can also detach ourselves and see them from afar. (Coffa 1991: 268, italics in original)

Why should this bear on the issue of semantic uniformity? Because the defender of truth in virtue of meaning can draw on this same distinction between how things appear 'from the inside', as it were, and how they really are when viewed from afar in order to resist the force of our intuitions regarding uniformity. From the inside, for the participants of a language, all statements of the language appear to be semantically uniform. They all seem to be about the world and answerable to it, including those which are analytic. This, the defender of truth in virtue of meaning can say, is only natural: these statements belong to the framework that shapes our thought about the world and so *from the inside* seem to be about the world because we conceive of the world as it presented by this framework. We are not immediately aware of the role these statements have in structuring our conception of the world, and so they are accorded equal status as the ordinary factual sentences of the language which they structure. However, upon reflection – considering the matter from

the outside, as it were – we are capable of recognising the fundamentally different roles these two kinds of statement play in the system, thereby allowing us to give up the illusion of semantic uniformity. Thus, when we construct our semantic theories like that of the disquotational schema, we do so from the inside perspective without any concern for the perspective outside the system. But in spite of this, it may nevertheless be that, viewed from outside the system, there are these two sets of statements with vastly different roles.

(iv) Concluding remarks

We began by observing our intuition that we typically suppose that we should accept a declarative sentence as true if and only if the world happens to be as it is represented by the sentence in question. We can now ask: how does the appeal to truth in virtue of meaning square with this intuition? In fact, from what we have seen, the question is harder to answer than we might have expected. On one hand, we might think that the very idea of truth in virtue of meaning clearly – for better or worse – violates the intuition. After all, it holds that there are sentences we are entitled to hold as true without any requirement that we are sensitive to how things are in the world; how the world is does not factor into our decision to endorse it as true. On the other hand, once we appreciate the different semantic roles the advocate of truth in virtue of meaning supposes analytic and synthetic statements to have, it is less clear that there is any violation of the intuition. If an analytic truth does not purport to say anything about the world, it does not demarcate any way for the world to be by which we should assess it – that is to say, it is not really a declarative sentence of the sort we initially took to be governed by the norm concerning our acceptance of sentences.

The proponent of truth in virtue of meaning is in a position to maintain that his strategy for dissolving the puzzle with which we started preserves the core of the initial intuition that our endorsement of a declarative sentence is to be guided by considerations of how the world is. He may grant that if a sentence represents the world as being some way or other, then it is true if and only if the world is that way, but may equally insist that we must be discerning as to whether, in any given instance, the antecedent of this

conditional is fulfilled. A sentence that does not have the role of genuinely representing the world is not open to assessment in the same way.

It should be clear by now that the suggestions we noted earlier by Arthur Sullivan and Gillian Russell, that attempts to explicate analyticity in terms of truth in virtue of meaning have been adversely affected by more recent developments in the philosophy of language, to be wide of the mark. For although there have been various clarifications and controversies about the philosophical conception of meaning over recent decades, there is little we have seen in the appeal to truth in virtue of meaning that leaves it hostage to the success or failure of any particular conception of meaning. For someone like Carnap, wherever there is a system of language, there are rules about what the expressions of the language are to mean, how they can be formed together into meaningful sentences and transformed by the substitution of one expression for another. So whatever the correct account of meaning, there will be a division to be drawn between those statements of the language that make claims about the objects of the domain and those statements are not themselves about the domain but rather articulate the rules which structure the language, and this division is all that is required for there to be some notion of truth in virtue of meaning.

Now, none of this is supposed to be a complete vindication of the truth in virtue of meaning strategy. The aim here has been rather more defensive, to show that the prevalent consensus that it has been decisively refuted, particularly by the Quinean criticisms of the notion, is misplaced. In the next chapters I consider more recent efforts to defend analyticity in a way that avoids appeal to anything like the idea of truth in virtue of meaning. One of the key issues brought to the fore by these more recent discussions of analyticity is that of how it could be possible for someone to understand an analytic truth and yet fail to know it; that is, if our knowledge of analytic truths is somehow to be rooted in our linguistic or conceptual practice (as, I suggested in the introduction, is what would make it analytic rather than rationally intuited), then how is it possible that we – the language users – could go wrong about them? This is a question to be faced by any account of analyticity, including that which appeals to some idea of truth in virtue of meaning. However, even if the considerations presented here do not amount

to a complete defence of the strategy, they nevertheless suggest that, if not in rude health, it is certainly in sufficiently robust shape as to indicate that the rumours of its demise have been greatly exaggerated.

The Epistemological Conception and Williamson's Challenge

In the previous chapter we examined the attempt to explicate analyticity in terms of truth in virtue of meaning, the idea that what accounts for the apriority of an analytic truth is that it is not really about the world in the same way as ordinary factual claims are. On this view, analytic truths articulate or embody the rules of a language that enable us to say things about the world but, in doing so, are neither about the world nor answerable to it, hence there is no mystery that we should be able to know them without having to investigate the world.

Criticisms of this strategy divide broadly into two kinds. Some, like Quine, are hostile to the very notion of analyticity *per se*, and oppose the idea of truth in virtue of meaning simply insofar as it is a kind of analyticity.¹ Others, however, are sympathetic to analyticity as an explanation of the a priori yet remain unconvinced by the Carnapian strategy. We find in Peacocke, for example, a dissatisfaction with the thought that different statements in a language could have distinct sources of truth, with one answering to how things are in the world and another to the rules of the language.² Part of his complaint is that, on the Carnapian view, the choice of rules for a language is a matter of convenience and nothing at all to do with truth.³ Perhaps more importantly, though, Peacocke also means to reject the Wittgensteinian conception of meaning that influenced Carnap. We noted in the previous chapter Wittgenstein's "picture" theory of meaning, according to which a proposition is about the world only if it is genuinely capable of being true or false, with the alleged consequence that one which is true in all possible situations does not represent a genuine possibility at all but rather serves to define "logical space." The same line of thought emerged in Carnap, for whom the fact that a sentence holds in all state-descriptions is indicative that its truth derives from the semantical rules of the language and not from the world. In contrast to this view, Peacocke writes

¹ See also BonJour 1998, particularly chapter 2.

² Peacocke 1993: 184

³ Ibid. 188

These claims of the *Tractatus* [that truth-functional tautologies say nothing and lack sense] are compulsory when it is also held both that ‘A proposition is a picture of reality’ (4.01) and that ‘There are no pictures that are true a priori’ (2.225). (...) When the support of the picture theory is removed, the early Wittgenstein’s reasons for saying that tautologies say nothing seem rather thin. He wrote that “A tautology has no truth-conditions, since it is unconditionally true” (4.461). What is unconditionally true is true; the tautology has truth-conditions, and they are always fulfilled. (Peacocke 1993: 183-4)

According to Peacocke, then, there are certain facts about the world which are necessary; the propositions that represent these facts have truth-conditions which are necessarily fulfilled. Yet, contra the Wittgensteinian view, that the truth-conditions are necessarily fulfilled does not entail that the proposition in question is not about the world; it is, rather, precisely because the proposition *is* about the world and the world necessarily fulfils those conditions.

Boghossian expresses a similar sentiment of dissatisfaction with the notion of truth in virtue of meaning when he says

In general, I have no idea what would constitute a better answer to the question: What is responsible for generating the truth of a given class of statements? than something bland like ‘the world’ or ‘the facts’. (Boghossian 1997: 336)

And again he writes,

Intuitively, the statements of logic appear to be fully factual statements, expressing objective truths about the world, even if necessary and (on occasion) obvious ones. (Boghossian 1997: 348)

So, if there are analytic truths (as Boghossian thinks there are), it is not by virtue of there being a set of statements which are somehow spared from answering to the world. Instead, it must be that what analytic truths there are derive their truth from how things are in the world. Boghossian thus proposes that we can separate what he perceives as the bad

metaphysical elements of analyticity, which are concerned with the issue of what makes an analytic statement true, from the respectable epistemological elements concerning how we could know them. In doing so, we can preserve that aspect which prompted our initial interest in analyticity without having to concede that there might be propositions which do not answer to the world for their truth.

The complaints expressed by Peacocke and Boghossian against truth in virtue of meaning are in fact of the same kind as those we saw in the previous chapter to be expressed by Williamson and Benaceraff regarding the strategy's rejection of semantic uniformity. We said then that the defender of truth in virtue of meaning has a comeback to offer against those, albeit that the strategy gives up the more intuitive stance with respect to the issue of semantic uniformity. The proposal that we can retain the epistemological elements of analyticity without the so-called "metaphysical" elements of truth in virtue of meaning is, in effect, an effort to secure a conception of analyticity that does not have to bear the cost of giving up our natural intuitions about semantic uniformity. The difficulty is that it reintroduces the problem of having to account for how there could be truths about the world, and thus answerable to it, which we are capable of knowing without having to investigate how things are in the world.

Both Peacocke and Boghossian defend accounts on which our knowledge of analytic truths is constitutive of our understanding them. It is not simply that our knowledge is *explained* by our understanding, knowledge of an analytic truth is built-in to the conditions for what it is to understand it; thus, understanding analytic truths is sufficient for us to know them. For Boghossian, the analyticity of certain conceptual truths, such as 'All vixens are female foxes', is explained by the fact that they involve the substitution of synonyms; to understand the constituent expressions is to know that they mean the same thing, thereby reducing the statement in question to a logical truth (for example, 'All vixens are female foxes' is equivalent to 'All vixens are vixens', which is an instance of the general logical truth $(x) (x = x)$). The apriority of logic is explained by the suggestion that certain statements implicitly define the logical constants, and so it is by holding these true that we give them the meanings they have. Anyone who understands the meaning of these logical expressions will thereby recognise that those statements which implicitly

define their semantic roles must be true. For example, Boghossian suggests, we might define 'if, then' by the following sentence:

If, if p then q, and p, then q.

The sentence itself uses the expressions 'if, then', but by stipulating that it is to be held true, we define the semantic roles of the expressions – their roles are those which would make the sentence itself true.⁴ And so anyone who grasps the meaning of 'if, then' knows that it denotes an inference pattern whose conditions of validity are satisfied by the sentence (and all other sentences of the same general form).

This talk of stipulating the truth of certain sentences in such a way as to give meaning to its constituent expressions can sound like it comes close to the idea that certain statements articulate rules which govern our use of expressions in a language, as on the truth in virtue of meaning strategy. Boghossian, however, insists that Wittgenstein and Carnap were in error to suppose that it follows from the fact that we stipulate the truth of certain sentences that they do not answer to the world for their truth. He appeals to Kripke's example of the introduction of 'meter', in which it is envisaged that someone introduces the expression by stipulating the truth of the sentence

Stick S is a meter long at t.

'Meter' then is defined in such a way as to denote the particular length had by stick S at t. But once it has been defined in this way, Boghossian argues, the sentence can be understood as making a particular claim about how things are in the world, namely, that the stick in question happens to have a certain length. That the stick is guaranteed to be of that length does not alter the fact, as he sees it, that its having that length is responsible for the truth of the sentence.⁵ Similarly, our stipulating that certain statements involving the logical constants are true does not make them any less dependent on the world for their truth. (There are serious difficulties with Boghossian's analogy here, in particular,

⁴ Boghossian 1997: 351

⁵ Ibid.

that whereas in the ‘meter’ example there is a quite demonstrable aspect of the world to which we can point in stipulating the truth of the sentence, it is less clear that this exists in the cases where we define the logical constants. This invites the worry that our stipulations are not responsive to anything in the world, thus making it difficult to see how they could be about the world in the sense Boghossian wants to maintain. In later developments of his idea, Boghossian attempts to ensure that his account can be truth-preserving by restricting it to a certain class of concepts from which our epistemic entitlements genuinely flow.⁶ These issues, though interesting, are extraneous to the present discussion and so are set aside here.)

Peacocke offers an account that is broadly similar to Boghossian’s in important respects.⁷ On his view, a concept is individuated according to some condition of what is required of a subject in order to possess the concept. This possession condition for a concept specifies a role within a subject’s “cognitive economy”, where fulfilling that role is constitutive of the concept in question. Since there is a particular role that a concept must play in order for a subject to count as possessing the concept, there are, accordingly, certain propositions which contain the concept that a subject must accept. That is, acceptance of those propositions is constitutive of a subject’s grasp of the relevant concept; one could not fail to accept those propositions without having failed to grasp some constituent concept. What ensures the truth of these propositions is that each possession condition is part of a relation where the other relatum is some entity that serves as the semantic content of the concept individuated by that particular possession condition. This relation between the possession condition for a concept and the entity that serves as its semantic content ensures that there is some feature of the world to which the concept – and the propositions that contain it – answers.

⁶ See Boghossian 2003a and 2003b.

⁷ Boghossian himself seems to acknowledge this when he says that he and Peacocke share a commitment to a Meaning-Entitlement Connection (MEC), according to which “any inferential transition built into the possession conditions for a concept are *eo ipso* prima facie entitling.” The main difference between their accounts, he suggests, is that whereas Boghossian tries to overcome the problems faced by MEC by restricting the thesis to a certain class of concepts (see top of the present page), Peacocke instead views it as applicable to all concepts but thinks we must therefore restrict what we count as a genuine concept. (Boghossian 2003b: 11, page number refers to the copy of this paper available on Paul Boghossian’s webpage.)

What is the relation between the Carnapian, metaphysical conception of analyticity and the epistemological conceptions defended by Peacocke and Boghossian? One could be forgiven for inferring from some of Boghossian's remarks that he intends the separation of their claims to be a schism rather than the evolution of a continued epistemological tradition. He writes,

I want to register my wholehearted agreement with Quine, that the metaphysical notion is of dubious explanatory value, and possibly also of dubious coherence. I believe that Quine's discrediting of this idea constitutes one of his most enduring contributions to philosophy. Fortunately for the analytic theory of the a priori, it can be shown that it need have nothing to do with the discredited idea. (Boghossian 1997: 335)

And again,

these two notions of analyticity are distinct, and...the analytic theory of the a priori needs only the epistemological notion and has no use whatsoever for the metaphysical one. (Boghossian 1997: 337)

However, it is important that the perceived need to draw the division does not obscure the commonalities between the two conceptions. On both, a subject's knowledge of an analytic truth derives from a sensitivity to the fact that its truth is a product of the rules governing the application of the constituent concepts or expressions. That is, analytic truths are those which are known by anyone who counts as understanding the relevant concepts or expressions. To fail to know an analytic truth, then, is to have in some way failed to grasp the meaning of at least one of the components. This heritage that the epistemological conception receives from its metaphysical predecessor is recognised by Peacocke, who argues that the common thread running through invocations of analyticity is the idea that there are certain propositions our acceptance of which is so intimately bound up with our grasp of them that we do not first entertain them and wonder as to their truth before choosing to accept them; rather, our decision to endorse them is to be explained simply by the fact that we grasp them. Referring to his own view, outlined above, as the 'metasemantic account', Peacocke writes,

There is one point on which the metasemantic theorist agrees with both Carnap and Quine, and indeed with the view of Wittgenstein from the 1930s onwards. A point of agreement between all these four views is their explicit or implicit rejection of a conception of meaning on which a thinker can first fully grasp the meaning of an expression and then, for any a priori proposition involving it, somehow or other move rationally from this prior grasp of meaning to the conclusion that the proposition is correct. The nature of this prior grasp is quite obscure, as is equally the nature of the alleged transition from it to an a priori proposition. It would have to be a transition which, in the nature of the case, cannot consist in any kind of inference. This point of agreement is one reason that the metasemantic account...should be seen as building on the thought of those grappling with these earlier problems in the 1930s, rather than as simply a reversion to some earlier view. (Peacocke 1993: 188)

The idea is that once a subject counts as having grasped an analytic truth, there is nothing more that can be done by way of acquiring further reason to endorse it. Crucially, there is no transitional or inferential step to be had between first grasping such a proposition and having the epistemic means by which to know it. Having grasped it, there is nothing more for a subject to do but endorse it. There is certainly something quite intuitive about this idea; our decision to endorse paradigmatic instances of analyticity seems to follow so immediately from entertaining them that it almost seems strange to think of it as a decision we take consciously at all. I do not, for example, first entertain the proposition that all vixens are female foxes and wonder as to its truth before then – as a separate step, as it were – deciding that it seems to be a reasonable enough claim to endorse. Instead, my acceptance of the proposition seems to be instantaneous with my grasping it precisely because its truth appears to me to be guaranteed by the meaning of the relevant expressions.

We might worry that this feature of some instances of analyticity does not extend, as Peacocke intends, to cover all those in which its advocates have been interested. The result of an even moderately complex mathematical proof, for example, would be claimed as analytic, and yet a subject's decision to affirm it may be entirely dependent on the

chain of reasoning by which it is derived from other mathematical truths.⁸ A subject may even question its truth, or whether it follows from the previous steps in the inferential process, so that even if the subject eventually has confidence in its truth, it would not be with the immediacy that is perhaps characteristic of the more paradigmatic instances of analyticity.

I am nevertheless inclined to think that there is something fundamentally right about the underlying point of Peacocke's analysis concerning invocations of analyticity, that they share the common notion that those instances in which a subject knows an analytic truth are ultimately rooted in the subject's grasp of the semantic role of the relevant expressions. In many cases, particularly those involving conceptual truths, there is no inferential or transitional step to be had from the subject's grasp of the proposition to the decision to endorse it; in such cases, we are sensitive to the fact that the proposition is one which articulates some aspect of the semantic role of the constituent concepts. Yet even where knowledge of an analytic truth is based on the kind of inferential process envisaged in the case of a mathematical proof, the process of deriving the truth can be seen to rest on the same non-inferential grasp of propositions as we find in the case of conceptual truths. For the inferential process can get off the ground only by presupposing the initial premises from which the result is derived as well as the rules of inference deployed in the process; on pain of regress, these cannot themselves be inferentially justified. For the proponent of analyticity, the warrant for the presupposition is supplied by the subject's grasp of the semantic roles of the relevant expressions. A subject who grasps the initial premises and the rules of inference may, by dint of this understanding, be capable of acquiring knowledge of the truths that can be derived from these. But to the extent that the inferential propositions are consequences of these more basic, non-inferential propositions, a subject's knowledge of these will ultimately be explained by the grasp of the semantic roles of the expressions involved.⁹

This idea, that a subject's knowledge of an analytic truth can be explained in terms of grasping the semantic roles of the relevant expressions, is what – for ease of presentation –

⁸ I owe this point to Quassim Cassam, who brought it to my attention in response to a paper presented at the MindGrad conference at the University of Warwick.

⁹ I am grateful to Marcus Giaquinto for comments that influenced my thinking about this point.

I shall call the *principle of epistemic sufficiency*. This is intended to capture the thought, common to the metaphysical and epistemological conceptions of analyticity, that the epistemic resources required to know an analytic truth are supplied entirely by a subject's grasp of the semantic roles of the relevant expressions. A subject who recognises a particular analytic statement as one that articulates some aspect of the semantic role of its constituent expressions is *ipso facto* entitled to endorse it; there is nothing more for the subject to do in order to be able to know it. If we are right to follow Peacocke in holding that something like this is a feature which defenders of analyticity have sought to respect as a central characteristic of their position, then it would seem that we could make the following claim: if there are such things as analytic truths, they must satisfy the principle of epistemic sufficiency; and, if there are no truths which can be said to satisfy the principle of epistemic sufficiency, there are no analytic truths.

Viewing matters this way can help illuminate the challenge posed by recent arguments against the notion of analyticity. Williamson in particular has argued repeatedly in recent writings that there are no truths our understanding of which could be sufficient to explain our knowledge of them since it is always possible for a subject to grasp a proposition – even a seemingly obvious one – and yet judge that it is not true. Mere grasp of a putatively analytic proposition cannot, then, be sufficient for knowing it, so in those cases where a subject does know such a truth, there must be something more that a subject does than simply grasp it. In this way, Williamson can be seen to challenge the principle of epistemic sufficiency: since there are no truths knowledge of which requires nothing more from a subject than mere understanding, there are none which satisfy the principle of epistemic sufficiency, and thus there are no analytic truths.

Williamson illustrates the point by considering what he takes to be a paradigmatic example of a putatively analytic truth, that all vixens are female foxes, presenting the cases of his deviant logicians Peter and Stephen. Peter has formed the considered view that the universal quantifier is existentially committing, believing that any sentence of the form 'All Fs are Gs' entails that there is at least one F. Being something of a sucker for conspiracy theories, he has also bought into the idea, propagated by a random website, that there are no such things as foxes, with all evidence to the contrary having been

planted by a government determined to make the morality of fox-hunting a political issue in order to deflect attention from concern about a potentially illegal invasion of another country. So Peter does not believe that all vixens are female foxes since, for him, to do so would commit him to the existence of foxes, which he rejects. Stephen, on the other hand, is a truth-value gap theorist who believes that borderline cases for vague terms produce circumstances in which sentences involving the vague term do not express determinate propositions capable of being true or false. Stephen also believes that there are female evolutionary ancestors of foxes which constitute borderline cases for application of the term 'vixen'. Since Stephen believes a sentence of the form 'All Fs are Gs' to be true only if the conditional 'x is F \rightarrow x is G' is true for all values of x, he believes the sentence 'All vixens are female foxes' to be false on the grounds that instances of the evolutionary ancestor constitute cases in which 'x is a vixen' is neither true nor false.

Do Peter and Stephen count as understanding the claim that all vixens are female foxes? It can be tempting to block Williamson's argument from the off by denying that they do; we might argue that there is a sense in which Peter, for example, in believing that the universal quantifier comes with existential import, is confused about the semantic role of our expressions 'all' and 'every'. We do not use them in that way, so to the extent that Peter believes that we do, he fails to understand the expressions and the sentences in which they feature. However, this response is a lot less appealing when we consider real-life analogues of the cases of Peter and Stephen, where there is genuine philosophical disagreement about the commitments of our logical vocabulary. Vann McGee, for instance, has argued that *modus ponens* is not truth preserving; Graham Priest has argued that one and the same proposition can be both true and false. Some of their opponents might want to insist that these views are simply the result of having failed to understand 'if' or 'not', but taking this line seems only to trivialize disagreement about substantive issues. If a subject genuinely misunderstands something, we can explain to them what we mean and help them correct their error; a misunderstanding is not the kind of thing in defence of which one can offer reasons. (One can, of course, offer an explanation – which we might call a reason – for how the misunderstanding arose, but that is not to defend the misunderstanding.) By contrast, McGee and Priest offer their conclusions as the results of considered philosophical arguments. It would not help resolve any disagreement we have

with them to simply reiterate what we mean by our uses of ‘if’ or ‘not’ – they know what we mean and they know what we suppose our commitments to be, yet having understood these things they nevertheless perceive there to be a problem with our view.¹⁰ To wave away their dissent, chalking it up to mere misunderstanding, therefore seems to miss the significance of their disagreement.

So it is, Williamson insists, in the cases of Peter and Stephen. In ordinary conversation they use the relevant expressions just as capably as anyone else; whatever idiosyncrasies they have as a result of their peculiar beliefs would remain indiscernible to the casual observer, unless conversation happens to impinge directly on topics relating to their pet theories in which case they are prone to sounding rather pedantic, though in this respect they are hardly different from many academics. They may be professors of philosophy, deferring to no-one in their logical abilities and having defended their views in the top journals. They understand what we mean by our use of the expressions and what we suppose our commitments to be – they just find there to be something problematic about the majority view. They are, Williamson says, no doubt wrong, but their errors are errors about logic and do not stem from any failure to grasp the proposition in question.

There can little doubt as to the significance of Williamson’s examples for attempts to defend the notion of analyticity. If there is no rejoinder to be had but to insist that Peter and Stephen simply do not understand the proposition that all vixens are female foxes, Williamson would have shown that analyticity requires a very narrow conception of what it is to count as understanding a proposition, one which lies well outside our ordinary attributions of understanding (for the reasons mentioned in connection with McGee and Priest) and so reduces considerably the explanatory virtues of claiming that our knowledge of a priori truths is explained by our understanding them. On the other hand, if the defender of analyticity concedes that Peter and Stephen understand the proposition, some account is required of how they could understand it and yet fail to know it. After all, if our knowledge of analytic truths is to be explained by grasp of the relevant concepts or

¹⁰ There are interesting questions to be addressed concerning what we are disagreeing about, but the present discussion does not require that we resolve these here. I defer offering an answer to these questions until chapter five.

expressions, then how is it possible that we – the users of the language – can go wrong about them?

The challenge is to identify an epistemically significant state that is common to all subjects who understand a given analytic truth but which allows for the satisfaction of the principle of epistemic sufficiency. If Peter and Stephen understand the proposition that all vixens are female foxes, any state they fail to be in with respect to the proposition is not one with which we may elucidate the notion of analyticity, since it would not be a state in which a subject is placed simply by virtue of understanding an analytic truth. But any state they *are* in would seem to undermine the principle of epistemic sufficiency, because although Peter and Stephen understand the proposition, they clearly do not know it. Consequently, it seems that something more than merely understanding the proposition must be required in order for a subject to know it, thus our knowledge cannot be explained by appeal to understanding alone and the principle of epistemic sufficiency is unsatisfied.

The problem generalises. It is always possible, for any candidate analytic truth, to construct cases like those of Peter and Stephen where a subject understands the relevant proposition but has some reason to refrain from endorsing it. Williamson takes this to show that there can be no truths our knowledge of which is explained by the fact that we understand them, since there must always be something more that a subject must do in addition to understanding a proposition in order to know it. So, on this view, there are no truths which satisfy the principle of epistemic sufficiency and hence no analytic truths.

However, the fact that a subject might understand an analytic truth but fail to know it does not tell against the possibility that in those cases where a subject *does* know the proposition, this knowledge is explained by the subject's understanding it. If we think about analyticity in terms of our entitlement to hold certain propositions as true just by virtue of understanding them, the fact that some people might not utilise their entitlement does not show that they had no entitlement to begin with. "They have blinded themselves," Williamson writes of Peter and Stephen, "but the upshot is still that they cannot see;" (Williamson 2006: 27) yet, to continue the metaphor, the light still

shines for the rest of us. Entitlements can be thrown away, but it does not mean they were never there. Nor does it mean they are thrown away at a cost to everyone else's as well.

What is required, then, is an account that preserves the idea that there are truths our knowledge of which is explained by our understanding them, while also making room for Williamson's point that it is possible for a subject to understand such a proposition and yet fail to know it. I want to suggest that such an account is afforded us by the notion of *being in a position to know*. For to speak of someone using the locution to 'being in a position to Φ ' is to indicate that the requisite resources for Φ -ing are available to that person, either in being already at their disposal or within reach (literal or metaphorical) in such a way that a reasonable amount of effort would equip them with those resources. Moreover, to say of someone that they are in a position to Φ leaves open the issue of whether they in fact Φ , since it is consistent with being in a position to Φ that one does not realise the potential to Φ at one's disposal.

Applied to the case of analyticity, the claim is that an analytic truth is one the grasp of which is sufficient to place a subject in a position to know it. If a subject's knowledge of a proposition were to be explained by the mere grasp of the proposition, then it would seem that anyone who likewise grasps it would thereby have the resources to know it at their disposal; that is, the entitlement utilised by the knowing subject would also be available to anyone else who grasped it. As the examples of Peter and Stephen highlight, this itself would not ensure that all those who grasp the proposition would know it. But we can say this: a subject in this situation who failed to know would not do so for want of any additional justification for the proposition; the justification available to them would be the same as that available to those who know it. And in having justification sufficient for knowledge at their disposal in this way, we can describe such a subject as being in a position to know the proposition, even though they in fact do not know it.

The suggestion immediately raises two important questions: what is it to be in a position to know in the sense that is relevant to our proposal?; and, how could understanding an analytic truth place a subject in a position to know it? A large part of the chapters that

follow is organised around providing answers to these questions. In this and the next two chapters I address the first; in chapter five I provide the basis for an answer to the second.

What has been said already provides a basic idea of what it is for a subject to be in a position to know in the relevant sense. A first approximation might be this: a subject S is in a position to know some proposition p if the justification for p at S 's disposal could be sufficient for someone to know that p . This is not as precise as it might be since there are issues surrounding what should count as a situation in which the justification is at a subject's disposal (I am inclined to think that the answers to such questions will vary between different modes of knowledge and so elude much further precision in this most general form). But irrespective of its defects, this initial formulation captures the intuitive idea that if there is a subject who knows that p on the basis of some justification, then if the same justification is at some other subject's disposal, then even if this latter subject does not know that p , there must nevertheless be a sense in which they are in a position to know that p . Moreover, the fact that a subject is in a position to know that p constrains the reasons for which p fails to be known. It cannot be because p is false because the notion of *being in a position to know* 'piggybacks', as it were, on the factivity of *knows*;¹¹ nor can it be for want of further justification on the part of the subject because, by hypothesis, the subject already has justification sufficient for knowledge at their disposal. Given this, the failure to know must be either by virtue of the subject being oblivious to the justification at their disposal, or – as in the cases of Peter and Stephen – by virtue of the subject having some false belief that undercuts the justification at their disposal. The idea of 'undercutting' here is that introduced by Pollock, according to which a proposition p is undercut by q in those cases where q does not directly rebut p but rather undermines the justification for p in such a way that were a subject to accept q , it would preclude utilising the justification for p .¹² In cases where a subject is in a position to know, so that the proposition is true and the justification to hand sufficient for knowledge, it seems that if the subject were to give up the undercutting belief – without acquiring any new one – we should expect them to utilise the justification at their disposal and thus know that p .

¹¹ For discussion of the factivity of knowledge, see Williamson 2000: 6 and 21.

¹² Pollock 1987: 38-9

We should be able to see, then, the difference between what it is for a subject to be in a position to know in the sense intended to be relevant to our discussion of analyticity and other occasions where we might find it natural to adopt the locution of 'being in a position to Φ ', or even other senses of 'being in a position to know'. The sense of the phrase as it applies to our discussion of analyticity is not the same as that I might use in claiming to be in a position to know where the Prime Minister will be tomorrow or to know which bands are playing locally at the weekend. Since many of the non-sensitive events in the Prime Minister's diary are a matter of public record, there may be a way of contacting his secretary or accessing a website in order to obtain the information. Similarly, I could pick up a local newspaper or enquire at the local concert hall to find out which bands are scheduled to make a performance over the coming days. In both of these cases there is a sense in which I am in a position to know insofar as there is something I can do in order to acquire the desired information. But this is not the sense in which we intend to speak of Peter and Stephen being in a position to know that all vixens are female foxes; there is nothing that Peter and Stephen need do in order to acquire the justification sufficient for knowledge of the proposition. The point of saying that they are in a position to know it is that they already have everything they need in order to know it.

Nor is the relevant sense of 'being in a position to know' akin to the sense in which it might be true of most adults that they are in a position to waggle their ears. Apparently, human bodies typically have muscles around the ears that would allow people to waggle their ears independently of other body parts; however, because these muscles are typically not exercised in the early years of a person's life, the ability to move one's ears in this way is often lost. Nonetheless, insofar as we have this capacity, we might suppose that there is a sense in which people are in a position to waggle their ears. And if so, then unlike my being in a position to know where the Prime Minister will be tomorrow, this would seem to be a case where everything needed is already at our disposal. Yet, still, the sense in which we might speak of 'being in a position to Φ ' does not seem to capture the situations of Peter and Stephen with respect to the proposition that all vixens are female foxes. The reason for this is that there is very little that many of us can do by the time we are adults about being able to move our ear muscles in this way. Even once a person is told that they are in a position to waggle their ears and the situation is explained to them, there is

generally nothing they can do to go from being merely in a position to move their ears to actually doing so. In this respect there is a significant difference with the cases of Peter and Stephen since there is something within their power that can be done in order for them to know that all vixens are female foxes, namely, give up their respective erroneous logical commitments that undercut their justification for the proposition. This is not to suggest that such a thing would be easy to do; they are, after all, supposed to be deeply entrenched commitments of people who have given a great deal of thought to their views. But the point is that it is something that is within their power to do – people who have thought hard about something can change their minds, however considered those views happen to be. The sense in which Peter and Stephen are in a position to know is one in which there is a genuine possibility that they could actualise that which they have the potential to do; were they to give up their peculiar logical commitments, then (provided they avoid some further pitfall) we should expect them to take advantage of their positive epistemic standing with respect to the proposition that all vixens are female foxes and know it.

These considerations motivate the idea that there does seem to be a distinct sense of ‘being in a position to know’ that not only represents the situations of Williamson’s deviant logicians but which is important in elucidating the role that understanding plays in our knowledge of analytic truths. A subject who understands an analytic truth is thereby in a position to know it in the sense that there is nothing more for the subject to do by way of acquiring the justification required in order to know it. Yet this leaves open the issue of whether the justification bestowed in understanding it is utilised by the subject or whether it faces, as in the cases of the deviant logicians, some sort of undercutting defeat that prevents the subject from knowing it.

Two issues confront the effort to defend analyticity in this way. First, Williamson himself foresees the appeal to *being in a position to know*, arguing that an account of analyticity elucidated in these terms remains susceptible to the problems presented by the cases of the deviant logicians. Second, it would be suspect if analyticity could be defended only by appeal to a notion whose application was itself limited to being deployed in defence of analyticity. And the distinct sense of ‘being in a position to know’, differing from other

ways we might use the locution of ‘being in a position to Φ ’, might increase the suspicion that some such special pleading is going on here. The next two chapters address these issues. The next chapter considers the idea of *being in a position to know* as it applies to cases of perceptual knowledge; I argue that these cases provide clear support for a similar sense of ‘being in a position to know’ as that required to elucidate analyticity. Moreover, I argue in the subsequent chapter, the parallel with cases of perceptual knowledge helps to illuminate the problem with Williamson’s attempt to resist the strategy of defending analyticity by appeal to *being in a position to know*.

Perceptual Knowledge and Being in a Position to Know

In this chapter I further explore the idea of *being in a position to know* introduced in response to Williamson's argument against the possibility of accommodating cases of disagreement about analytic truths. In particular, I argue that far from being an *ad hoc* manoeuvre deployed solely in defence of analyticity, the same sense of *being in a position to know* has a role to play in describing certain features of our perceptual knowledge. In section (i) I sketch an initial motivation for this view drawing on an exchange between Brewer and Martin in which the latter attributes to *being in a position to know* a place within McDowell's account of perceptual knowledge. In section (ii) I consider a potential problem for applying the notion in perceptual cases, observing that fake barn-style scenarios look like paradigmatic instances where it should apply. If, however, subjects in these cases do not have knowledge (as many suppose) then neither are they in a position to know. This might suggest that there is an inherent instability in the idea of *being in a position to know*.

I argue that the apparent problem arises from the competing influences of fallibilist and infallibilist approaches to perceptual entitlement. Neither of these by themselves is at odds with the notion of being in a position to know, but we need to be careful about separating their respective commitments. In section (iii) I outline an infallibilist approach to perceptual entitlement as advocated by McDowell, arguing that there is evidence within McDowell's own writing that he views subjects in fake barn cases as not only being in a position to know but also as actually knowing the proposition in question. In section (iv) I outline a fallibilist account of entitlement, arguing that although there is room on this view for a subject being in a position to know, the fallibilist can maintain that the conditions for being so are not realised in fake barn-style cases; so subjects in these scenarios neither know nor are they in a position to know. In this way, I argue, there is a place to be had for *being in a position to know* on our accounts of perceptual knowledge, irrespective of whether we endorse fallibilism or infallibilism about perceptual entitlement.

(i) Accommodating perceptual defeasibility

We typically suppose that our perceptual experiences provide us with knowledge about the world. Such experiences represent to us the world's being some particular way; in many instances we take our experience at face value, accepting the world to be as it is represented to us. In believing that at least some of these instances yield knowledge of the world, we suppose that experience is capable of supplying a warrant for our judgement that things are as they appear. Successful cases of perception, in which the world is as it is represented and we are appropriately related to the world's being that way, provide us with an entitlement to the beliefs we form about the world when we take experience at face value. Although there may be occasions when circumstances conspire as to prevent a subject from utilising the warrant provided by even a veridical experience, in favourable circumstances the entitlements supplied by these experiences enable the beliefs we form in response to be instances of knowledge.

We may note that cases of perceptual knowledge appear to conform to an analogue of the principle of epistemic sufficiency we saw in the previous chapter to be characteristic of analyticity. When a subject has a perceptual experience, there is nothing more to be done by way of acquiring further reason to endorse its content. Nothing about the experience as it is presented to the subject will guarantee that it is not deceptive, so those cases in which a subject knows that the world is as it is represented require that the subject take things at face value. In this respect, there is a clear asymmetry with those cases where a subject fails to know on the basis of a veridical experience. Suppose I am presented with a tiled wall which appears to me to be blue, and suppose further that it is. I might be inclined to judge on the basis of my experience that the tiles before me are blue before noticing that everything in the room has a blue hue because the room is lit by a fluorescent light above. In such a scenario I would have a reason to abstain from my initial inclination to judge that the tiles before me are blue. Moreover, there would in this situation be a transitional step from first noticing the way things are presented in experience to then deciding not to endorse appearances, namely, noticing that there is a fluorescent light tinting the whole room, which undermines my initial reason for supposing the wall tiles to be blue. However, in cases where the subject endorses appearances, there is no positive evidence that can be adduced on the basis of the

experience that would establish its veridicality. Nor does there seem to be any inferential route from the mere phenomenology of experience to the knowledge we ordinarily suppose perceptual experience to provide.¹ Instances of perceptual knowledge, then, have no transitional step by which a subject can acquire warrant for the move from grasping the content of an experience to accepting that things are as they appear. Rather, the warrant for a subject's judgement stems directly from the experience itself; a subject who in favourable circumstances takes at face value how things are presented by a veridical experience is entitled to endorse that things are as they appear.

With this picture of our perceptual experience in the foreground, one question to arise is: why should the experience of the world's being a certain way entitle us to endorse that it is so? Brewer, in his *Perception and Reason*, offers one proposal of how we might answer this question: that a subject's very grasp of the content presented by a veridical experience is sufficient to warrant endorsing that content. Starting from the idea that, in perceptual experience, objects are presented as determinately located relative to the subject-centric perspective of the perceiver, Brewer argues that the identifying knowledge necessary for referring to particular objects is provided by the subject's standing in certain "perceptual-attentional" relations to those objects. Possessing a sufficiently discriminating conception of an object involves exploiting the way experience displays the location of the object in space; since no two objects (at least of the same sort) can occupy the very same spatio-temporal location, the location in which an object is presented relative to the subject-centric perspective provides a means of uniquely identifying the object and thus of being able to entertain singular thoughts about it. A subject's perceptual judgement of some object *a* that *a* is *F* is, on this account, a judgement that *that* object – presented to the subject as being *there* – is *F*. In this way, Brewer argues, judgements formed on the basis of perceptual experience contain an irreducible demonstrative component.

That the content of a perceptual demonstrative of this sort should be recognised by a subject to be a particular mind-independent object is, Brewer suggests, the upshot of a relational egocentric identification of a location "in a world of places and things which are quite independent of [the subject's] actual experiences of them." (Brewer 1999: 190) This

¹ On this point, see McDowell 1998c: 396-399

enables the subject to appreciate that how things are presented in experience is jointly dependent on how the objects in question are distributed in the world around him together with his contemporaneous location amongst them, thereby allowing the subject to make sense of the possibility that “the thing in question, *at just that location*, might have been displayed as differently located – relative to him – from elsewhere.” (Brewer 1999: 197, italics in original) This, Brewer supposes, is crucial because

successful reference to a persisting *mind-independent* thing depends on more than there happening to be such a thing at the end of his pointing finger when he thinks ‘that is F’, say. He must have some appreciation of the fact that his thought is the joint upshot of the way that thing is and his meeting some further, independent enabling conditions upon recognising this – being in the right place, looking in the right direction, with sufficient illumination and so on. For it is precisely his grasp of the possibility that these further, independent conditions might fail to obtain which enables him to make sense of the possibility that that very thing might have been just as it is without his recognising it; and this, in turn, is essential to his understanding of the demonstrative by which he identifies the relevant object *as making reference to a mind-independent thing* whose existence is quite independent of his awareness of it. (Brewer 2000: 430, italics in original)

Why should we care about the subject’s awareness of the objects of perceptual reference as things which are mind-independent? The significance, according to Brewer, lies in a subject’s appreciation that things as they are presented in experience might have been otherwise, either because things as they are in the world might have been different or because the subject might not have stood in the particular perceptual-attentional relation to the objects of experience. In appreciating this, the subject is able to recognise that the content of an experience – that *that (there)* is F – is one that could be entertained only if things are as the experience presents them as being, thereby providing the entitlement for the subject to judge that they are so. Brewer writes,

For, simply in virtue of entertaining perceptual demonstrative contents of this kind, he recognises that it is that thing, there in relation to him, say, if determinacy of singular reference is secured by grasp of spatial location, and mind-independently

thus, which is currently displayed.... That is to say, he understands that his current apprehension that things are thus and so is in part due to the very fact that they are. His grasping the content that that is thus is in part due to the fact that that *is* thus. He therefore recognises the relevant content *as* his apprehension of the facts, his *epistemic openness* to the way things mind-independently are out there. (Brewer 1999: 204, italics in original)

There are, as Brewer himself notes, parallels between his account of perceptual knowledge and the accounts of analyticity offered by the likes of Peacocke and Boghossian.² Both are guided by the effort to preserve the seeming immediacy with which a subject is able to take advantage of the warrant supplied by their respective sources of knowledge. Both, as a result of this, propose that we view a subject's knowledge of a proposition as *constitutive* of grasping it, so that a subject's knowledge is simply a consequence of grasping the proposition. We saw in the previous chapter that, in the case of analyticity, this proposal is susceptible to the objection raised by Williamson that a subject might understand an analytic truth and yet fail to know it. And Brewer's account of perceptual knowledge is open to a similar line of response.

Martin points out that the subject of a veridical perceptual experience might grasp its content and nevertheless refrain from endorsing that things are as they appear to be. He illustrates the point by considering a case in which a machine is developed capable of causing people to see perfect hallucinations of oranges, perhaps brought out from time to time by a professor to demonstrate the phenomenal indistinguishability of veridical and hallucinatory perceptual experience. Suppose that a subject who knows of the machine and is aware of its capabilities agrees to undergo the experience. However, unbeknown to the subject, the machine is broken and unable to cause any hallucinations. So if it appears to the subject that he is seeing an orange, it can in these circumstances only be because there is an orange in front of him and he is seeing it. In such a scenario the subject may still entertain the proposition that there is an orange in front of him, and because the experience is a veridical one there is no question but that any demonstrative component of the thought would successfully refer. Yet despite grasping this content, it is not only

² Brewer 2000: 430-31

natural but also appropriate that the subject in this situation should refrain from endorsing that things are in fact as they are presented as being in experience. That is, in addition to the experiential content of seeing that there is an orange in front of him, the subject also has reason to think that the judgement he would typically be inclined to make is one to which he may not be entitled since he considers it a live possibility that he is a victim of the machine's deceptive capabilities. The presence of such a doubt renders the subject unable to utilise whatever warrant the experience, in normal circumstances, might have provided for judging that there is an orange in front of him.

Martin's point is not (as Williamson's with analyticity) that there is therefore no knowledge of the relevant sort to be had. His point is rather that there must be a looser relation between the grasp of a perceptual content and the knowledge of it than on Brewer's account: knowledge of a perceptual content cannot be constitutive of a subject's grasping it, even in veridical cases, since a subject might grasp such a content and yet fail to know it; so even if we want to hold on to the idea that a subject's perceptual experience can be sufficient to warrant the beliefs formed in response, we must nevertheless make room for this possibility.

Martin suggests that Brewer is able to accommodate defeasibility of this sort within the general account of perceptual knowledge he otherwise wants to offer. Central to this account, he observes, is the idea that what entitles a subject to the perceptual judgement that things are as they represented in experience is the fact that it is formed on the basis of his "epistemic openness" to the world, as displayed by the object-dependence of the demonstrative component of a perceptual judgement that *that (there)* is thus. Behind this lies an infallibilist conception of warrant, according to which a subject's experience in cases of veridical perception provides warrant for a judgement that is absent in cases of illusion or hallucination. According to Brewer, a subject has warrant for the judgement that *that (there)* is thus because it is a content that could be entertained by the subject only if he is appropriately related to the objects in question, that is, only if things are as they appear to him. There are, notes Martin, two corollaries to such a conception of warrant: first, a subject who is the unwitting victim of an illusion would suppose that he is warranted in the judgement he is inclined to make even though, in such a case, he is not;

secondly, for a subject to have defeating reasons for the warrant a veridical experience otherwise supplies is for him to think it possible that he is in a situation with no reason at all for the demonstrative judgement he is inclined to make.³ He goes on to write,

Following John McDowell, who endorses both of these features of perceptual justification,...one might simply claim that perceiving that something is the case puts one in a position to know that that is so. What is distinctive of perception is that it is an openness to the world, that how things are in one's environment make themselves manifest to one. Having things so manifest puts one in a privileged epistemic position such that one can then come to have knowledge about how one's environment is. There are situations – cases of illusion or hallucination – when one is not so epistemically privileged; the world fails to be manifest to one, even though from the inside, reflecting on one's situation, one would not be able to tell that the situation is so. Someone who is alive to a genuine possibility of illusion or hallucination will allow doubts about whether they really are perceiving to constrain the judgements he or she makes. Even if a subject is in the privileged situation, the presence of doubts can put him or her in a position where he or she cannot responsibly exploit the privilege that they have. And it is in this way that collateral information can defeat the justificatory role that otherwise perception would play. (Martin 2001: 447)

Martin's invocation of the idea of *being in a position to know* here is slightly elusive since he neither goes on to explicate the sense of the phrase he takes to be relevant (even though, as we noted in the previous chapter, there are various ways we find it natural to speak of someone being in a position to Φ), nor is it made clear precisely how the McDowellian features of justification should lend support to the idea (the phrase 'being in a position to know' is not one that, so far as I can tell, appears in the locus of McDowell's epistemological writings).⁴ But Martin's remarks do suggest the idea found in McDowell of veridical experience as that which presents a subject with, as it were, a direct contact with some fact about the world, which in turn makes the knowledge of that fact available to the subject.⁵ The privileged position in which a subject is placed by this relation to the

³ Martin 2001: 447

⁴ The obvious sources here are his (1998c), (1998d) and (1998e). It has been pointed out to me by Craig French that McDowell has used the phrase in work written subsequent to Martin's remark in his (2008).

⁵ See McDowell 1998d: 386-7 and 1998e: 417

world can be exploited so long as the subject is not hindered by having reason to think that he is not entitled to take appearances at face value. McDowell's clearest statement of this general view is in his 'Knowledge by Hearsay' in discussing knowledge acquired through testimony, his account of which parallels his treatment of perceptual knowledge. He writes that "if a knower gives intelligible expression to his knowledge, he puts it into the public domain, where it can be picked up by those who understand the expression, as long as the opportunity is not closed to them because it would be doxastically irresponsible to believe the speaker." (McDowell 1998e: 438) The idea is that when someone articulates something they know, a fact about the world is made directly available to those who hear it, enabling those entitled to take things at face value to know the fact in question.⁶ And it is the parallel view of perceptual knowledge that Martin supposes to be amenable to the kind of defeasibility relevant to the case of the subject who worries about being the victim of a machine-induced hallucination.

I take it that the point of Martin's suggestion that we might understand perceiving something to be the case as putting one in a position to know it is that, on the McDowellian view he presents, although a subject's experience may be one which directly relates him to some fact about the world, the mere availability of the fact is no guarantee that the subject will be able to exploit the privileged position in which he stands. If a subject has reason to suppose that things as they are presented should not be taken at face value, then it will not be appropriate to judge on the basis of experience that things are as they appear. Yet a subject in this situation might nevertheless stand in the same warrant-conferring relation to the world as someone who *is* capable of exploiting their position and of thus knowing how things are on the basis of their experience. In this sense, the respective situations of these subjects might both be characterised as one of

⁶ McDowell's only qualification to this picture is to point out that this does not mean that doxastic responsibility in addition to things being as they are presented are jointly sufficient for knowledge, since there may be cases in which an apparent perception is not a genuine perception for lack of the appropriate relation between the fact and the experience. (McDowell 1998e: 431-2) I assume that the kinds of case he has in mind are those such as veridical hallucination, where an object might be causally responsible for a hallucinatory experience that happens to represent things as they actually are. In these kinds of cases, there may be nothing about the constraints of doxastic responsibility that should prevent a subject from taking things at face value, but they would not be instances of knowledge given that it is not the fact about the world to which the subject's judgement is responsive. There may be parallel cases involving testimony in which there is no doxastic irresponsibility in accepting someone's word, and what is said is true but not an expression of knowledge.

being in a position to know, even though only one in fact knows while one does not. Moreover, we might suppose that were we to correct the false, defeating belief that undermined the subject's ability to take things at face value, the subject would then be able to judge on the basis of experience that things are as they appear. For example, were someone to inform the subject worried by the prospect of being the victim of a machine-induced hallucination that the machine is broken and incapable of creating any experience as of an orange in front of him, then so long as he did not immediately acquire some new defeating belief, he should come to know on the basis of his experience that there is an orange in front of him.

(ii) Fake barn-style cases

It should be clear, then, that the idea of *being in a position to know* that Martin takes to be relevant to the accounts of Brewer and McDowell mirrors in crucial respects the notion we suggested in the previous chapter to be relevant to analyticity. So we have *prima facie* evidence to suggest that it is an idea with genuine epistemological import independent of its deployment in defence of analyticity. However, we may yet worry that there is an instability inherent in the notion that threatens to undermine the independent support offered by cases of perceptual knowledge, since we seem to be able to run an argument to show that the sense of 'being in a position to know' discussed up to this point runs contrary to a powerful and common intuition that proved influential in discussions of perceptual and testimonial knowledge. To see this, let us consider the following cases.

*Fake Barns*⁷

Henry is driving through the countryside, paying careful attention to the scenery as he goes along. Each object he sees is fully in view and exhibits features characteristic of its type. Henry sees what appears to him to be a barn and judges that it is so. He is right; the object in question is a barn. Unbeknown to Henry, however, he happens to be driving through an area where there are numerous fake barns which are indistinguishable in appearance from the real barn – models that are good imitations but unable to be used as barns. Henry just happens to have picked out the one real barn in his view.

⁷ The example is from Goldman 1988, though often attributed to Carl Ginet.

*Censored Newspapers*⁸

A businessman arrives from abroad at an airport in a country governed by a dictatorship. He goes to the newsstand and picks up the nearest newspaper to hand, believing the information he reads. Unbeknown to him, although all the newspaper journalists in the country are reliable and honest, the government regularly censors and revises the content of the country's newspapers. However, the department responsible for carrying out these revisions is under-staffed and is able to amend some, but not all, of the newspapers on any given day. Those that go unrevised invariably print the truth. On this particular day, the newspaper the businessman picks up has evaded revision.

These cases are designed to elicit the intuition that although Henry and the businessman each form true beliefs on the basis of the respective perceptual and testimonial experiences, the beliefs they form are not instances of knowledge. The underlying thought in these cases is similar to our reasons for supposing that a subject who has reason to think he might be the victim of a machine-induced hallucination is thereby unable to judge that things are as they appear, even when his is a genuine perceptual experience. Faced with the disjunction that either things are as they appear or the experience is merely hallucinatory, a subject should recognise that he lacks the ability to discriminate between the two possibilities to be able to judge that it is one rather than the other that obtains. Even if the subject were to flagrantly disregard the epistemic peril of the situation by deciding to press on and believe that things are as they appear, we would not count this as an instance of knowledge on the grounds that it would be merely a matter of luck that he happened to opt for the true disjunct rather than the false one when he had no reason to choose between them.

Similarly, although neither Henry nor the businessman are aware of the pitfalls in their respective situations, their true beliefs have the same flavour of luck as the subject who is alive to the possibility of deception. They too lack the ability to discriminate between the veridical experience they each happen to have and a deceptive experience which in the described circumstances represents a very close possibility. Were they to have employed the same belief-forming method in the close possible worlds where they alight on the

⁸ The example is from Peacocke 1986.

‘wrong’ object (seeing one of the fake barns in Henry’s case, picking up a different, revised newspaper in the businessman’s), the resultant beliefs formed in response to those experiences would have been false. So, again, we have the same sense that the true beliefs they happen to form are merely accidental in a way that precludes them being instances of knowledge. Or, to make the point in a slightly different way, it seems clear that if Henry or the businessman were to be made aware of the presence of fake barns in the vicinity or that the government is prone to altering the content of its national newspapers, we would naturally expect them to refrain from judging that things are as they appear. We should expect that a person in their situation would realise that they could not possibly tell whether the content of their experience is veridical or whether it is the result of some deception. And as with the previous case, even if they were to carry on regardless and judge things to be as they appear, we would not count this belief as knowledge since they would have no basis that one or other situation obtains. Given that if they were to be aware of their circumstances, they would not know that things are as they appear, it seems counterintuitive to suppose that they could know when they lack this additional information; as Jason Stanley puts it, “It does not seem correct that adding a little ignorance increases knowledge.” (Stanley 2005: 7)⁹

We might express the general thought that unifies the responses to these cases in following way:

(K) If the experience on which basis a subject forms a belief leaves open that things are other than as they appear, so that the subject would have the same belief in close possible worlds in which it is false, then the subject’s belief cannot be knowledge.¹⁰

This seems to handle both kinds of the relevant cases. In those where the subject is aware that there is a real possibility of deception, there should be the accompanying realisation

⁹ Stanley’s remark is made in the context of speaking about the way in which practical interest may have a role in whether it is appropriate to attribute knowledge to a subject, rather than being made in direct reference to fake barn-style cases, but his point is an intuitive one that seems to have equal application here.

¹⁰ This is a close variation on Peacocke’s slightly modified version of the third condition in Nozick’s analysis of knowledge. Peacocke’s claim is that “a belief is knowledge only if it is reached by a total method such that use of any of its proper or improper initial submethods yields in the actual and in nearby possible worlds a belief true in the world in which it is used.” (Peacocke 1986: 134)

that the judgement he is instinctively inclined to make on the basis of his experience is one he would also be inclined to make if he were being deceived. The subject's inclination to judge that things are a certain way is therefore not responsive to how they in fact are, since he lacks the discriminatory capacity to discern on the basis of the experience whether things are as they appear or whether he is deceived, hence it is doxastically irresponsible for him to judge that things are as they appear. In fake barn-style scenarios, the subject is not aware of the possibility of deception, yet we still suppose that if it is merely a matter of luck that he forms the true judgement on the basis of an experience that could just as easily have lead him astray, then the belief is not one so appropriately related to the fact as to be knowledge.

Here, then, is the problem we might raise for the appeal to *being in a position to know* in cases of perceptual knowledge. The significance of the notion lay in the ability to preserve the idea that a subject who has a veridical perceptual experience is entitled to judge that things are a certain way, so that, if circumstances are favourable, he can know that it is so. Of equal importance to the notion is the idea that a subject who has the same warrant-conferring experience as one who knows, yet finding himself in unfavourable circumstances that prevent him from knowing the content of the experience, can nevertheless be described as being in a position to know it in the sense that he has an experience that in other circumstances would allow him to know its content. From this, we might suppose that if Henry has the veridical experience of seeing a barn, he should either know that there is a barn or, at the very least, count as being in a position to know that there is a barn, since his experience is one that in favourable circumstances would allow him to know that there is a barn. But it seems that if Henry fails to know that there is a barn then neither can he be said to be in a position to know that there is a barn. For the notion of *being in a position to know* shares its normative features – the epistemic conditions a subject must meet – with those of our more commonplace notion of *knowledge*, while leaving open whether or not a subject utilises the position he is in by virtue of meeting those conditions. So, if the kind of luck enjoyed by Henry in fake barn country is deemed incompatible with his knowing that there is a barn, then luck of this sort would seem also to be incompatible with him being in a position to know that there is a barn. And this appears to be borne out when we consider the idea of 'correcting' the

subject's faulty belief. We have said that if a subject is in a position to know, but does not know due to some undercutting, false belief, then if that belief were corrected, the subject would – all other things being equal – know that which he was in a position to know. In Henry's case, however, the relevant false belief is that he is in a normal perceptual environment where there are numerous barns in the vicinity. As we have already noted, if Henry were to be made aware of the details concerning his situation, he would not thereby be able to know that he sees a barn; rather, considerations of doxastic responsibility should prevent him from judging that he sees a barn. In this respect also, Henry's epistemic situation does not seem to be one of being in a position to know.

So, on the one hand, Henry's situation looks like it is ideally suited for the idea of *being in a position to know* insofar as he has a veridical experience that, in different circumstances, would enable him to know that things are as they appear to him. On the other hand, we have reason to think that if Henry does not know that there is a barn then neither is he in a position to know that there is a barn. We might worry that this seeming tension is the product of an instability inherent in the idea of *being in a position to know*, at least as it might apply in cases of perceptual knowledge, though even that qualification would bring with it a dialectical concern. For the purpose of showing that *being in a position to know* has application in cases of perceptual knowledge was to provide independent support for drawing on the notion in defending analyticity from Williamson's objection. The envisaged concession concerning the stability of the notion in cases of perceptual knowledge would remove any such support, perhaps giving rise to the worry that the move in defence of analyticity is undermotivated or *ad hoc*, leaving us ill-placed for responding to Williamson.

I want to suggest that the tension brought out by the foregoing concern arises not out of any internal instability in the notion of *being in a position to know* but rather out of conflating two much broader approaches to the issue of perceptual entitlement, specifically an *infallibilist* conception of entitlement according to which veridical experience provides a subject with a warrant that is absent in non-veridical experience and a *fallibilist* conception according to which the warrant available to a subject in

veridical experience may also be available in non-veridical experience.¹¹ Advocates of these different approaches have adopted very different stances over what to say about fake barn-style cases, and if we are not careful to distinguish their competing influences, we may be liable to diagnose the resultant tensions as inherent to the idea of *being in a position to know*. My aim here is not to adjudicate between these views, but rather to show that *being in a position to know* is precluded by neither. On an infallibilist conception of perceptual entitlement, there is reason to think that the subject in fake barn country is not only in a position to know that there is barn but even – as I think McDowell is inclined to claim – that he may know that there is a barn, whereas on a fallibilist conception the conditions for being in a position to know simply fail to be realised in fake barn-style cases. So, while both can accommodate the idea of *being in a position to know*, they have different conditions for when a subject counts as being in such a state, delivering conflicting verdicts about what should be said in tricky cases such as that presented by the subject in fake barn country.

(iii) Infallibilism about perceptual entitlement

Let us start with the infallibilist view, the view that the warrant bestowed upon a subject who has a veridical experience is one that would not be similarly available if the subject were having a perfectly matching non-veridical experience. First of all, what reason would anyone have for believing that? For someone like McDowell – whom I shall take as representative of the view – the motivation is bound up with the effort to secure a conception of knowledge capable of preserving the idea that it is something fundamentally different from merely accidental true belief, which in turn is crucial if we are to preserve the many ordinary things we take ourselves to know about the world in the face of sceptical challenges to those beliefs. Scepticism gets its foothold by pressing the case that experience provides no reason to judge that things are as they appear to us to be rather than as the upshot of some widespread deception: ‘How do you know that there’s

¹¹ The labels ‘fallibilism’ and ‘infallibilism’ are sometimes used in quite different ways by epistemologists. Fallibilism, as I have introduced it into the present discussion, is to be distinguished from (though not necessarily wholly unrelated to) the idea that any of our beliefs about the world could be false, and *mutatis mutandis* for infallibilism and the idea that there are truths about the world we can know with certainty beyond reproach. I am unaware of there being many detailed discussions of the distinction, as I have drawn it, in the literature, though Pryor (2000) and Wright (2002) provide some, while McDowell (1998c, 1998d, 1998e) and Burge (2003) offer paradigmatic versions of the infallibilist and fallibilist positions respectively. My own appreciation of the distinction owes much to conversations about the topic with Mike Martin and Marcus Giaquinto.

an external world?', asks the sceptic; 'Because I see that I have a hand, which entails the existence of an external world', we might reply; 'But surely one could have that same experience if one were deceived by an evil demon', the sceptic responds, 'so how can you rule out that possibility in such a way as to have reason for your belief that there is an external world?' The worry is that if we have nothing to offer by way of response to this question, then even if our beliefs happen to be true – there is an external world and we do see our hands – it would be only a matter of good fortune rather than because experience itself provides us with conclusive reason to think them true.

Now, McDowell does not suggest that we should feel imperilled by a sceptical threat of this kind; nevertheless, he argues that non-sceptical epistemologies on which the world itself is not a constitutive feature of veridical experience leave themselves open to precisely the problem to which the sceptic points. He argues that if the world is not a constitutive feature of veridical experience, then for all a subject knows on the basis of experience, there is always the possibility that things are other than as they appear. By hypothesis, the "experiential intake" is the same in veridical and non-veridical cases, so if non-veridical experience does not relate a subject to some fact about the world, then neither is a subject related to any fact about the world in veridical experience. Thus, McDowell writes, "when this capacity [to tell by looking, on the basis of one's experiential intake] does yield knowledge, we have to conceive the basis as a *highest common factor* of what is available to experience in the deceptive and non-deceptive cases alike, and hence as something that is at best a defeasible ground for the knowledge." (McDowell 1998d: 386, italics in original) The problem with this, he argues, is that since for all that is presented to the subject in experience, things could be other than as they appear. Even if we could establish somehow that, over the course of experience, it is on any given occasion more likely that we are not deceived than that we are, this would still leave the truth of any particular judgement made on the basis of experience a hostage to fortune in a way we would not typically suppose knowledge to be. McDowell offers a scenario in which someone plays the roulette wheel: assuming the game is not fixed, we would not suppose that the person who correctly calls the colour on which the ball will land knew that it would land that way since it could equally have gone the other way; yet even if there were a thousand slots, nine hundred and ninety-nine one colour and only one of the

other, we should still refrain from supposing that the roulette player knew what the outcome would be, since for all he knew before the wheel was spun, things might have defied the odds and turned out otherwise.¹² And so it is, McDowell argues, with perceptual knowledge: if a subject makes a judgement on the basis of an experience that leaves open the possibility that things could be other than as they appear, then even if correct it is a matter of luck that is thereby precluded from being knowledge.¹³

McDowell's proposal is that we should understand perceptual experience as a direct openness to the world, where the world itself is constitutive of a subject's experience. (This is the idea to which we saw Martin appeal in recommending what he took to be the McDowellian notion of *being in a position to know* in order to accommodate the defeasibility supposed to be problematic for Brewer.) The experience had by a subject when in this direct relation to the world is one that would be unavailable were things to be other than as they appear, since in a case of the latter sort the world itself is not a feature of the experience. On this view, then, McDowell writes,

statements of such forms [as 'I see that...'] are proper moves in the game of giving reasons, and their truth fully vindicates entitlement to the embedded propositions. (...) Someone who can truly make a claim of that form has an entitlement, *incompatible with any possibility of falsehood*, to a claim whose content is given by the embedded proposition. The entitlement consists in the visual availability to her of the fact she would affirm in making that claim. (McDowell 2002: 98, italics added)

So whereas the sceptical worry gains traction if we grant that one may have the same experience in perceptual and deceptive cases alike, a subject who can truly assert 'I see that I have hands' has conclusive reason to reject the overtures of the sceptic, since the evidence cited is an experience that would not be available were the sceptical hypothesis to be realised, namely, the experience of seeing that there is a hand.¹⁴ Such a judgement is

¹² McDowell 1998e: 422

¹³ McDowell 1998c: 403-405

¹⁴ There are interesting issues here about whether the subject in this situation could really be entitled to draw on evidence of this kind to rule out the sceptical possibilities, or whether the warrant the subject has for believing that there is a hand fails to transmit to a conclusion of this kind (see Wright (2002) for a discussion of this issue as it relates directly to McDowell's account). Since the main purpose of the present discussion is to show that the notion of being in a position to know is compatible with

the result of a subject's openness to the world and the absence of any defeating considerations that prevent the subject from exploiting the privileged position in which he is placed by this direct relation to the world.

The relevance of all this to fake barn-style cases is the bearing it has on the issue of what it is for a subject to be lucky in a way that precludes knowledge. McDowell distinguishes between two different kinds of luck a perceiver might enjoy. There is, on the one hand, the rather basic sort of luck we have whenever the world turns out to be as it is presented in experience. Since there is no way by which we can establish whether a particular experience is veridical or otherwise, we are always hostage to the world being presented to us other than as it really is; McDowell writes, "we are vulnerable to the world's playing us false; and when the world does not play us false we are indebted to it." (1998c: 407) On the other hand, there is the luck one enjoys when, for all one knows on the basis of the evidence at one's disposal, things might be other than as one supposes, yet one nevertheless happens to be right. This is the luck exhibited by the roulette player who calls the colour correctly in advance of the spin; for all he knows, the ball could just as easily have landed on a different colour. The difference between these two kinds of luck is significant: the former is not only compatible with knowledge but is an inescapable element of knowledge – we are always reliant on the world doing us the favour of being the way it appears;¹⁵ the latter, by contrast, is incompatible with knowledge because the evidence on which basis the judgement is formed is not sufficient to determine that things will turn out to be one way rather than another.

With this distinction in mind, we can revisit the case of Henry in fake barn country. The prevalent response, which the case is designed to elicit, is that Henry does not know on the basis of his experience that there is a barn because – so the thought goes – it is a matter of luck that his gaze happens to alight on the one genuine barn in the vicinity when he forms his judgement. But we can now see that this is problematic not if he is lucky *per se* but only if the luck he enjoys is of certain kind that is incompatible with

infallibilist and fallibilist accounts, rather than showing either of these to be right, it is not integral that these issues are pursued further here. That said, I am not inclined to think that there is a problem for the McDowellian view for reasons much along the same lines as those offered by Beebe in her (2001).

¹⁵ McDowell 1998c: 405-408

knowledge. In this case, however, the common intuition is that the luck is of this kind. We earlier articulated the principle behind the intuitions in cases of this sort as

(K) If the experience on which basis a subject forms a belief leaves open that things are other than as they appear, so that the subject would have the same belief in close possible worlds in which it is false, then the subject's belief cannot be knowledge.

The problem for Henry, then, is that there seem to be very close possible worlds in which his gaze happens to alight upon one of the numerous fake barns in the area that would have been indistinguishable for him from the genuine barn he happens to see. Since Henry lacks the ability to discriminate between the genuine barns and the fakes, then he would have made a judgement of the same type if his experience had been of a fake barn. In such a scenario, his judgement would be false; so if there is nothing about his epistemic situation that differs from this other than the fact that his judgement happens to be true – that is, if for all he knows on the basis of his experience, things might just have turned out otherwise – then his getting hold of the truth is little more than an accident of the kind that we typically regard as incompatible with knowledge.

Yet, on McDowell's view, it would be mistaken to suppose that Henry is lucky in this way, for in having alighted upon the genuine barn, Henry's judgement is made on the basis of an experience that would not have been available to him had he been looking at a fake barn. Henry's experience is one of standing in a genuine perceptual relation to an object of a certain kind and seeing the world as it is; standing in this relation to the world affords Henry an entitlement – “incompatible with any possibility of falsehood” – to judge that things are as they are presented in experience. If Henry were to have looked upon one of the nearby fakes, it would appear to him as though there is a barn even though there is not; no fact of the relevant kind would be made available to him in a way that would entitle him to judge that things are as they appear. So in those close possible worlds, Henry would judge that there is a barn, but – crucially, for McDowell – he would not do so on the same experiential basis as when he has a genuine perceptual experience of a barn.

We noted earlier McDowell's view that knowledge is made available to a subject if they are related by experience to some fact about the world, a position they are capable of exploiting in the absence of defeating considerations that would make it doxastically irresponsible for them to judge that things are as they appear. Since Henry's is a genuine perceptual experience in circumstances where it is not doxastically irresponsible for him to take appearances at face value, then if he is not adjudged to be lucky in a way that precludes knowledge, it seems that he should know that there is a barn. This is in keeping with McDowell's verdict about a structurally identical case involving testimonial knowledge.¹⁶ In a twist on the tale of the boy who cries 'Wolf' until the townspeople no longer believe him, McDowell considers the situation of a stranger who, having never visited the locale before, happens to hear the boy's cry the one time there really is a wolf nearby. The scenario is the same as fake barn-style cases in the relevant respects. There are numerous other cries of 'Wolf' which the stranger might equally have heard, though only the cry he happens to hear is truthful. The townspeople who have heard the boy's previous cries and found them to be false are in a position where it is impossible for them to responsibly believe that the boy is alerting them to the presence of a wolf, though the stranger has no knowledge of these previous falsehoods. And, of course, he lacks any information or capacity that would enable him to discriminate between a true or false cry from the boy. It is therefore reasonable to suppose that had the stranger been present on any of the previous occasions, he would have been equally inclined to believe that there is a wolf nearby. Does the stranger know on the basis of his testimonial experience that there is a wolf? McDowell's verdict is that he does; he writes,

The apparatus I am recommending allows us to entertain the idea that the stranger might acquire knowledge from the boy.... I do not find such a possibility obviously offensive to intuition. The case would be one in which something that might otherwise be an opportunity for the acquisition of knowledge is closed to those who know too much. (McDowell 1998e: 436)

In this case, McDowell considers the epistemic situation of the townspeople to be akin to that of the subject in Martin's example who is wary of being deceived by a machine-

¹⁶ I am grateful to Craig French, who reminded me of this significant passage and of its relevance for the McDowellian view in which I am interested here.

induced hallucination. They, like him, have an experience which presents them with a fact about the world, viz. the presence of a wolf, which in normal circumstances would enable them to know that this fact obtains. However, they have information regarding the reliability of the boy – just as the subject in Martin’s example has information regarding the capabilities of the machine – that makes it doxastically irresponsible for them to believe that things are as their experience presents them to be, thus rendering them unable to exploit the privileged position in which their experience places them. But to the extent that they stand in this relation to the relevant fact, they are in a position to know it. By contrast, neither the stranger in McDowell’s example nor Henry in fake barn country have information concerning the potential for error in the judgements they are respectively inclined to make. They likewise stand in a relation to some fact about the world. Their ignorance of those factors that might otherwise defeat their belief enables them to exploit the position in which they are placed by their experiences.

The conception of knowledge around which the foregoing account is centred is not one averse to the idea that, as Stanley put it, “a little ignorance increases knowledge.” For knowledge acquired through experience is not the product of our gathering as much information as we can about our situation, piecing the various bits together and forming the beliefs that achieve the best fit. It is rather the result of being able to exploit the position in which we are placed when we are directly related to a fact about the world by experience. Our ability to accept that we are so related is partly a matter of doxastic responsibility. There may be circumstances when we are related to some fact in this way yet the possession of certain information about our situation undermines our ability to exploit our epistemic position, although the absence of this information would have left the ability to exploit our position intact.

(iv) Fallibilism about perceptual entitlement

It would perhaps be inaccurate to suggest that anyone inclined towards infallibilism about perceptual entitlement need be committed to following this McDowellian view about fake barn-style cases. There is, we may suppose, room for someone to subscribe to the infallibilist thesis that a subject’s veridical experience provides an entitlement that would be unavailable in the case of a perfectly matching non-veridical experience while also

maintaining that the entitlement bestowed by a veridical experience can, in unfavourable circumstances, be undermined in such a way as to render the subject unable to know that things are as they are presented in experience. Such a view would deliver a similar verdict to that offered by the fallibilist about perceptual entitlement about fake barn cases. The fallibilist holds that the entitlement bestowed upon a subject by a veridical experience may also be available in non-veridical cases. This is perhaps the more common view to be found in the literature, even if rarely aired explicitly as a thesis.¹⁷ For someone like Burge, who does articulate and defend the view, fallibilism follows from the idea that the warrant for a perceptual belief derives from the well-functioning of the system by which it is produced. Burge argues that an epistemic good, such as warrant, is to be understood and assessed relative to its contribution to achieving the fundamental good of any representational system, namely, truth. So the perceptual beliefs generated by a system are warranted if the system itself is a good route by which one may acquire knowledge of the environment. This imposes a condition of reliability that a perceptual system must meet if the beliefs it generates are to be warranted: such a system must be reliable in accurately representing the environment in which it normally operates in a way that generates true beliefs about that environment. A subject whose perceptual system meets this condition is thereby entitled to the beliefs formed on the basis of his experience. Burge writes,

it is built into the nature of the competence associated with the formation of a reliable perceptual state that the state make a non-accidental, explanatory, positive contribution to true belief and knowledge.... Being in a reliable perceptual state in itself makes a positive contribution to a believer's having a prima facie entitlement to form an appropriately conceptualized perceptual belief from it. (Burge 2003: 530)

On this view, then, perceptual entitlement is compatible with the possibility of falsehood. A subject may be entitled to a particular perceptual belief by virtue of having a reliable perceptual system prone to generating true beliefs in its normal operating environment even though, on this particular occasion, things are not as they appear. Perceptual entitlement is compatible with what Burge calls “brute error – error that does not reflect

¹⁷ Cf. Pryor 2000: 518

badly on the functioning or use of the individual's representational systems, relative to the function of representing veridically." (Burge 2003: 507) Entitlement is also compatible with "brute truth," that is, cases where the generation of a true belief is accidental relative to the good use and functioning of a subject's perceptual system. Such cases are not instances of knowledge since, by hypothesis, the subject's use of the perceptual system is not a reliable route to veridicality, thus making the truth of beliefs generated in these cases a matter of luck in a way that precludes their being knowledge.

Fake barn-style cases would seem to be of this latter sort. In fake barn country, Henry's presumably otherwise reliable perceptual system is removed from its normal operating environment, since normally there are not indistinguishable fakes of any variety in the vicinity. In this new environment his perceptual system is no longer a reliable route to true belief, at least concerning the presence of barns. Were Henry's gaze to have alighted upon one of the fake barns, his experience would have been of the same sort as when he sees the genuine barn, on which basis he is inclined to judge that there is barn. Since in this environment his perceptual system is liable to mislead him, generating false beliefs rather than true ones, it is a matter of good fortune and not by any virtue of the means by which Henry's perceptual beliefs are formed that he happens to judge truly that there is a barn. Moreover, even taking into account the different kinds of luck we noted in connection with McDowell's view, the good fortune here would seem to be of a kind that is incompatible with knowledge. For in contrast with someone like McDowell, the fallibilist treats veridical and non-veridical experience as of the same fundamental kind, hence the fact that a subject can be entitled to the perceptual beliefs formed on the basis of either. So, on a fallibilist account of the situation, for all Henry knows on the basis of his experience, he could have equally been looking at a fake barn. So there are very close possible worlds in which Henry would have the same experience yet in which his judgement would be false.

Henry therefore does not know there is a barn. Nor, then, is he in a position to know that there is a barn. For the point of saying that someone is in a position to know, in the sense relevant to our discussion here, is that the subject has at his disposal the epistemic resources required for knowledge of a proposition while leaving open the issue of whether

the subject takes advantage of this. According to the fallibilist, the possibility of knowledge that there is a barn is closed off to Henry due to the contextual unreliability of his perceptual system, making it a matter of luck that he has the true belief that there is a barn. If we accept this verdict, then Henry cannot be said to be in a position to know in this situation either since there is a crucial epistemic resource required for knowledge which he lacks, viz. a contextually reliable perceptual system that makes a non-accidental contribution to the judgements he is inclined to make in that environment. The presence of fake barns undermines the possibility that a subject could come to know on the basis of perceptual experience that there is a barn; and if the possibility of knowledge is closed off in this way, then a subject in such circumstances cannot be in a position to know either.

This is not to suggest, of course, that there is no room on a fallibilist conception of perceptual entitlement for the notion of *being in a position to know*. But in the same way that the fallibilist disagrees with the infallibilist stance regarding the source and nature of perceptual entitlement, so will this be mirrored by a disagreement about the conditions required for a subject to be in a position to know. In many cases the differences between them may be immaterial to the verdict they deliver. For both, the subject in Martin's example has the resources required for knowledge at his disposal, though their explanations for this diverge – for the infallibilist it is explained by the subject's openness to the world in having an experience that would not be available were he not seeing the world as it is; for the fallibilist it is explained by the subject's veridical experience that results from a reliable perceptual system. And both can accommodate the idea that the subject's belief concerning the possibility of machine-induced hallucination undermines the warrant otherwise available to be exploited. Thus, for both there is a clear sense in which the subject can be said to be in a position to know. However, the differing conceptions of perceptual entitlement may in certain cases lead to conflicting verdicts being delivered as to whether a subject in a given situation should count as being in a position to know, and this is what we see in fake barn-style cases.

(v) Concluding remarks

We can now see that the competing intuitions concerning what we should say about fake barn cases arise not from any tension inherent in the notion of *being in a position to*

know but rather due to the pressures of these two different conceptions of perceptual entitlement. It is perhaps tempting, given the infallibilist picture in connection with which Martin discusses the idea of *being in a position to know*, to suppose that the notion ought to have application in fake barn-style cases where the subject has a veridical experience of a kind an infallibilist may take to be incompatible with the luck identified by (K). And for all that has been argued here, this may be right. But it is not a stance that is imposed upon us by the notion of *being in a position to know per se*. For we have seen that *being in a position to know* can equally be accommodated within a fallibilist approach to perceptual entitlement which is capable of explaining why the notion need have no application in fake barn cases.

The upshot is that *being in a position to know* appears to signify an epistemic state that has application with respect to features of our perceptual knowledge, irrespective of whether we are inclined to adopt a fallibilist or infallibilist stance about perceptual entitlement. This offers some support for the suggestion that we can draw upon the notion in the effort to elucidate analyticity in a way that addresses the challenge we saw to be laid down by Williamson in the previous chapter. The support, however, is defeasible: admitting a place for being in a position to know in an account of our perceptual knowledge does not entail that one must accept that it is capable of underwriting an account of analyticity. Williamson argues that an account based along these lines is equally afflicted by the challenge of accommodating the kind of disagreement he envisages. I turn to these arguments in the next chapter.

Analyticity, Understanding and Being in a Position to Know

We have seen that Williamson rejects an understanding-based account of our knowledge of truths such as *All vixens are female foxes* on the grounds that it is always possible to understand what such a claim expresses and yet not know it. Since there is this possibility for any candidate analytic truth that a subject may understand it and not know it, Williamson argues, it cannot be that those who do know it do so by virtue of understanding alone. However, as I have already argued, it does not follow from the fact that one might understand an analytic truth and not know it that in those cases where a subject does know it, that knowledge should not be explained in terms of the subject's understanding. In the previous two chapters I have suggested that the notion of *being in a position to know* offers a viable description of an epistemic state to which we can appeal in elucidating the phenomenon of analyticity. This enables us to accommodate the kind of disagreement that Williamson envisages while preserving the idea that there can be truths our knowledge of which is explained by our understanding them.

A move of this sort is anticipated by Williamson, who thinks it unable to underwrite an account of analyticity that is not equally susceptible to the difficulties introduced by cases in which a subject counts as genuinely understanding a putatively analytic truth without having any inclination to endorse it. His remarks have at times suggested that the problem lies with the idea of *being in a position to know*, though the line of criticism he pursues appears to target the wrong sense of this phrase. Across the breadth of his writing on analyticity, however, Williamson's primary target is the conception of understanding he takes to be required in order to support an account of analyticity. The challenge in this respect is to find some picture of understanding that, on one hand, is sufficiently weak as to allow that one can always understand an analytic truth without endorsing it while, on the other hand, is sufficiently strong to explain our knowledge of analytic truths. Williamson's claim is that a conception of understanding that is weak enough to accommodate cases like those of his deviant logicians is too weak to be capable of the epistemic work required for analyticity.

In this chapter I examine the arguments offered by Williamson to support this pessimism about the idea that understanding an analytic truth could place a subject in a position to know it. In section (i) I consider an objection he raises against the viability of the appeal to *being in a position to know*. I argue that by drawing on the parallel with perceptual knowledge in the previous chapter, we can see that the argument misidentifies the sense of the phrase relevant to the defence of analyticity. In section (ii) I present Williamson's concern that understanding cannot do the epistemic work analyticity requires, outlining the conception of understanding he supposes we must accept in order to accommodate cases like those of the deviant logicians. According to this view, all there is to understanding a language is the ability to interact with other speakers; since this presupposes nothing about what a subject must accept in order to count as understanding, there is no truth a subject could count as knowing just by virtue of understanding it. In (iii) I argue that the ability to interact with others, while important, does not exhaust the available senses of what it is to count as understanding. Understanding an utterance involves grasping the mutual significance of the expressions used by speakers in a linguistic community, thereby knowing which thought is expressed by a speaker's utterance using those expressions. This offers a stronger conception of understanding than that presented by Williamson. In (iv) I show that people like the deviant logicians still count as understanding even on this stronger account. Thus, we can resist Williamson's claim that the need to accommodate cases of disagreement forces us to accept the weak conception of understanding he presents.

(i) Distinguishing the psychological and the epistemic

In his paper 'Conceptual Truth', where Williamson initially considers the response that understanding an analytic truth might put a subject in a position to know it, Williamson writes that

The intended point of the phrase 'in a position to' here is that the knowledge is readily available, not merely that some non-observational psychological route of unspecified complexity leads ultimately to the knowledge in question, for the latter reading [would be] too weak to be of much interest. (Williamson 2006: 26)

It is perhaps not entirely clear what Williamson envisages by the latter option here, though I take it that he is, at least in part, concerned to rule out the idea that if there exists an a priori argument establishing some truth, the mere existence of such an argument is not enough for one to be in a position to know its conclusion; to say of someone that they are in a position to know a given proposition is to indicate a much closer connection than that it is merely possible that they *could* come to know it by some means – it indicates a sense in which the necessary epistemic work has already been done, making knowledge of the proposition (to re-use Williamson’s phrase) “readily available”. But this phrase introduces its own difficulties. The sense Williamson supposes to be relevant seems to be the same he has in mind in *Knowledge and Its Limits*, where he writes: “If one is in a position to know p , and one has done what one is in a position to do to decide whether p is true, then one does know p . The fact is open to one’s view, unhidden even if one does not yet see it.” (Williamson 2000: 95) In this sense, the knowledge that all vixens are female foxes is not readily available to Williamson’s deviant logicians Peter and Stephen. They have done what they are in a position to do to decide whether the claim is true. Yet their deviant logical beliefs are so deeply entrenched that, psychologically, they are incapable of giving up their respective commitments. Moreover, Williamson argues, even if they were capable of pressing ahead and believing the proposition, they would be behaving in a way so deeply irrational to themselves that their doubts would surely defeat whatever warrant their belief might otherwise have had. Accordingly, Peter and Stephen fail not only to know the proposition but also to be in a position to know it.

The problem is that Williamson’s argument appears to turn on a conception of what it is to be in a position to know which differs from that previously identified as relevant to the defence of analyticity. For even when we accept the constraint that to be in a position to know the knowledge must in some sense be readily available, we can draw a distinction between a subject’s possession of the epistemic resources required in order to know some proposition and that subject’s psychological capability of exploiting those resources in such a way as to realise their potential. Williamson’s argument clearly supposes that the ready availability required to be in a position to know is one on which these epistemic and psychological aspects coincide. Yet, from what we have seen in the previous chapters,

there is room for a conception of *being in a position to know* on which a subject might have the epistemic resources necessary for knowledge of a given proposition while nevertheless being psychologically incapable of exploiting that privileged epistemic position. It is helpful here to consider again the analogy with cases of perceptual knowledge. The subject who has a veridical perceptual experience of an object but has reason to worry that the experience is a machine-induced hallucination has the epistemic resources required to know that there is an object of a particular sort before him. Like the deviant logicians, if the subject were to disregard the threat of deception and press ahead with the belief that things are as they appear, he would be behaving in such a way as to defeat whatever warrant his experience might otherwise have been capable of providing. Thus, in this situation the subject cannot exploit the position in which he is placed by the experience. Nevertheless, we may incline towards the view that the subject's possession of epistemic resources that would (but for his belief about the heightened possibility of deception) allow him to know the proposition in question makes it appropriate to describe his situation as one of being in a position to know.

If this is right, the relevant sense of *being in a position to know* is one on which the epistemic aspect concerning the warrant at the subject's disposal and the psychological aspect concerning the subject's ability to exploit his epistemic position do not coincide: the subject has the epistemic resources necessary for knowledge but, in the circumstances in which he finds himself, lacks the capability to exploit them. And so it is, the defender of analyticity can insist, in the cases of Williamson's deviant logicians. Insofar as their understanding of the proposition that all vixens are female foxes provides the epistemic resources necessary for knowing it, thereby making the knowledge "readily available", Peter and Stephen are in a position to know it. However, due to the erroneous logical commitments with which they have allowed themselves to become saddled, they are rendered incapable of taking advantage of their epistemic position.

The issue here is not merely a terminological dispute about how we ought to use the phrase 'being in a position to know'. Williamson's argument requires that there be no sense in which it is appropriate to describe the epistemic state in which we are placed by understanding a putatively analytic truth as one of being in a position to know; that is,

whatever viable conception of *being in a position to know* we might propose, Williamson claims that the deviant logicians fail to be in that state. The problem is that Williamson's argument for this claim rests on the fact that, from a psychological perspective, Peter and Stephen are incapable of relinquishing the commitments that leave them unable to know that all vixens are female foxes. We can see, however, that if we are careful to distinguish between the epistemic and the psychological aspects concerning a subject's being in a position to know, we can make sense of the idea that a subject might have the epistemic resources but be unable to exploit them. So the fact that a subject may be incapable of utilising the epistemic resources available to them fails to show that the subject is not in a position to know.

(ii) Williamson's conception of understanding

In *The Philosophy of Philosophy*, Williamson attempts to block the kind of move made in the previous section in defence of analyticity, arguing that the picture of understanding required to support analyticity is untenable. At the core of his argument is the idea that once we make room for cases like those of the deviant logicians – conceding that, for any truth, one could understand it and yet fail to know it – we are committed to a conception of understanding so weak as to be incapable of doing the epistemic work necessary for analyticity.

Williamson considers a defence of analyticity that construes the relevant knowledge-understanding connection in the following way:

(KU) Whoever knows "Every vixen is a vixen" in the normal way does so simply on the basis of their understanding of the sentence.¹

A formulation along these lines would seem to be amenable to the suggestion pursued in the preceding chapters that understanding an analytic truth puts a subject in a position to know it, since it posits a subject's understanding as the relevant explanation for the

¹ This construal of the relevant knowledge-understanding connection is presented by Williamson as a weaker alternative to the stronger connections that represent the stances of proponents of analyticity like Boghossian or Peacocke, such as "Necessarily, whoever understands the sentence "Every vixen is a female fox" knows "Every vixen is a female fox"." (Williamson 2007: 77) or even "Necessarily, whoever understands the sentence "Every vixen is a vixen" has a disposition to assent to it" (Ibid: 100).

knowledge in question while leaving open the possibility that a subject might understand the sentence and yet not know it. But, Williamson argues, despite its initial plausibility, there is no conception of understanding available that would serve to make (KU) true. We can highlight the problem, he suggests, by reflecting on what is denoted by the definite description “their understanding of the sentence”, dividing possible interpretations into thick and thin candidates. The thin candidates are the mere fact that they understand the sentence and so are exactly the same for any two people who count as understanding it. The thick candidates are the underlying facts that constitute the respective thin candidates, “the facts that realize this particular subject’s understanding at this particular time” (Williamson 2007: 131), and as such may differ between any two people who understand the sentence since there may be different underlying facts that constitute their respective states of understanding.

By dividing the candidates in this way, Williamson argues, we can see that neither yields a conception of understanding suitable as a basis for analyticity. Suppose “their understanding of the sentence denotes the thick candidates. These are the facts that constitute a particular subject’s understanding of a sentence, including various cognitive capacities that are not in general necessary for understanding the sentence, but which serve to make up that subject’s competence with it. So, Williamson says, the facts that constitute Peter’s understanding of ‘Every vixen is a female fox’ include his logical capacities, while Stephen’s understanding of the same sentence is constituted by facts which include his rather different logical capacities. But since the various cognitive capacities we might cite from person to person will involve things that are not required *per se* to understand the sentence, it is not clear how we could expect to offer a general explanation of our knowledge in terms of these. The thick candidates are thus too thick to do the work required for analyticity, since they involve cognitive capacities that are not semantic in any relevant sense.

Apparently, the thin candidates do not fare any better. These are the mere fact that a subject understands a sentence and, Williamson claims, imply no specific logical capacity at all as Peter and Stephen, with their different (and erroneous) logical commitments, demonstrate. If a subject can understand the sentence and not know it, then, so the

argument goes, there is no particular capacity that is presupposed by a subject's semantic competence with a sentence. And since there is nothing presupposed by the fact that a subject is semantically competent with a sentence, the mere fact that a subject understands a sentence is by itself insufficient to place a subject in any particular epistemic state with respect to that proposition. Thus, the thin candidates present too thin a conception of understanding to provide the kind of knowledge-understanding connection supposed by (KU).

Williamson concludes that "the attempt to base the epistemology of obvious truths such as ['Every vixen is a vixen' or 'All vixens are female foxes'] on preconditions for understanding them rests on a false conception of understanding." (Williamson 2007: 133) This naturally raises a question about what kind of conception of understanding Williamson would have us endorse and, moreover, why this conception should be viewed as so unfavourable to the notion of analyticity. His answer to the former is constructed around the need to have conditions for understanding sufficiently weak to accommodate cases like those of Peter and Stephen. The fullest statements are presented in the following two passages:

Understanding words in a natural language has much to do with the ability to use them in ways that facilitate smooth and fruitful interaction with other members of the community. That ability can be realized in indefinitely various forms. Speakers can compensate for their deviance on one point by their orthodoxy on others, their ability to predict the reactions of non-deviant speakers, their willingness in the long run to have their utterances evaluated by public standards. (...) [S]uch compensation is often possible when the deviance results from localized interference in the normal practice of using a word by high-level theoretical concerns. Thus there is no litmus test for understanding. Whatever local test is proposed, someone could fail it and still do well enough elsewhere with the word to count as understanding it. (Williamson 2007: 97)

Each individual uses words as words of a public language; their meanings are constitutively determined not individually but socially, through the spectrum of linguistic activity across the community as a whole. The social determination of

meaning requires nothing like an exact match in use between different individuals; it requires only enough connection in use between them to form a social practice. Full participation in that practice constitutes full understanding. That is why there is no litmus test for understanding. (Williamson 2007: 98)

The picture Williamson offers here is hardly so unusual as to load the argument in his favour. We may note that the view finds a parallel in that defended by Wiggins in his paper 'Languages as Social Objects', where he argues that "a language like English or Polish is a social object, a public thing with attributes irreducible to the individual psychology of speakers." (Wiggins 1997: 499) There, Wiggins appears to endorse a similar conception of understanding to that offered by Williamson, writing that "Knowing a language is simply being able to hear well enough, to interpret well enough and to speak well enough and aptly enough some sufficiency of the sentences that have currency or acceptance in that language." (Wiggins 1997: 508) The emphasis here is on the ability to interact successfully with other speakers of the language; understanding is achieved so long as a subject is able to hear and speak the language "well enough" to facilitate communication. So a subject may deviate from other speakers of the language in the use of an expression and still count as understanding, provided this deviancy does not interfere with the overarching goal of interactive success.

In fact, we may construe a line of reasoning by Gillian Russell as suggesting that the conception of understanding presented by Williamson is actually a natural consequence of the major work in the philosophy of language in recent decades. She notes that it has been a major theme of post-Quinean philosophy of language that relatively little is demanded of us in order to count as semantically competent.² Putnam's work on natural kind expressions and Kripke's on proper names share the idea that what it takes to be competent with these expressions is that one use them in such a way as to defer to the communal practice, intending to use them as other members of the linguistic community use them.³ Little is demanded by way of knowledge or discriminating information concerning the extension of an expression in order to count as competent with it. A subject may even have a number of false beliefs about an object, such as to cause him to

² Russell 2010: 190

³ Putnam (1975), Kripke (1980)

withhold assent from certain basic truths about it, and yet nevertheless count as being competent with the expression by which he refers to it.

This, as we have seen, poses a problem for a view of analyticity on which a subject's assent is supposed to be a precondition for understanding a sentence. For if the conditions for understanding are sufficiently weak, it can no longer be claimed that a subject's failure to assent to a seemingly obvious truth – perhaps known by most familiar with the constituent expressions – must be the result of a failure to understand it. Russell articulates the threat neatly, writing,

The ideas of reference via causal-historical chains, division of linguistic labour, and semantic deference provide models of language on which users of an expression need not have knowledge of anything (far less a meaning) that would uniquely determine the extension of the expression, and so they tend to allow for very minimal requirements on semantic competence.... (...) This lowering of the standards for semantic competence threatens to erode the epistemic import of analyticity. If all analytic truths can be known on the basis of what you have to know to be semantically competent, then...the less you have to know to be semantically competent, then fewer truths will be analytic. (Russell 2010: 190)

One way for us to understand Williamson's concern about analyticity, then, is as the idea that if we take seriously this lowering of the standards for semantic competence and follow it through to its natural conclusion, there can be no truths our knowledge of which can be explained merely in terms of our competence with the relevant expressions. For if there is nothing more to understanding than the ability to interact with fellow participants in a linguistic community, so that counting as competent is consistent with holding various false beliefs about the meaning of an expression, then for any given expression there is no particular claim to which a subject need assent in order to count as understanding that expression. That is, for any expression e and any sentence $S(e)$ in which it features, it is always possible that a subject may dissent from $S(e)$ and yet count as understanding e in virtue of his more general ability to use the expression in interacting with other speakers. Consequently, there can be no sentence assent to which is licensed

by understanding alone, since one could just as easily withhold assent and still equally count as understanding it.

Understanding cannot do the epistemic work required by analyticity because, on Williamson's view, it is not a matter of one's preparedness to assent to any particular sentence or set of sentences – it is much more holistic than that, determined by one's overall practice of using an expression to interact with other speakers. This weakens the conditions for understanding to a point that it becomes difficult to see how mere understanding alone could explain a subject's acceptance of any given truth, since dissent is quite compatible with a subject's competence. So it is not, as the above quotation from Russell suggests, that there are simply fewer analytic truths than we might have supposed: there are none.

(iii) Understanding and knowledge

The problem for the defender of analyticity is that there is much that is attractive about the picture Williamson presents. It is a central fact about language that it is a means by which we communicate with others, so it seems natural that we might evaluate someone's linguistic competence by the success with which they communicate, and it is a merit of Williamson's account that this feature is given such a central place in the view he outlines. Moreover, assuming – as we suggested – that Williamson's view is a natural extension of externalist developments in the philosophy of language over recent decades, it would be disappointing for the defender of analyticity to find that this otherwise attractive picture had to be denied in order to resist Williamson's objections against analyticity.

A better response, I want to suggest, is to accept that Williamson articulates a perfectly good sense in which we might want to count people as understanding while denying that this sense exhausts those available to which the defender of analyticity can appeal. Williamson's argument, as with that concerning *being in a position to know* in section (i), turns on there being no sense of what it is to understand a language that can do the work required for analyticity; if we can show this not to be so, we leave open the possibility that a defence of analyticity can be provided. What we require, then, is some conception of understanding stronger than that invoked by Williamson, capable of underwriting an

account of analyticity. Yet, ideally, we should also desire that even on this stronger account of understanding, people like the deviant logicians still count as meeting the standards for semantic competence. The reason for this is the same as that offered for initially granting to Williamson that they understand (see p. 61). Examples like these do not appear to be simple cases of misunderstanding, particularly when we consider real-life analogues involving disagreement about logic: it just doesn't ring true to say, for example, that Vann McGee fails to understand what we mean by 'if' or Graham Priest what we mean by 'not'. Their views may be incorrect, but they certainly know what we take ourselves to mean by the relevant expressions, and it would not enlighten them to simply reiterate our stance in contrast to theirs. The challenge, therefore, is to identify some stronger conception of understanding that is nevertheless capable of accommodating Williamson's insight that there is always the possibility that, for any truth, a subject might understand without assenting to it.

We can begin to motivate a different picture to that offered by Williamson by reflecting on the mutual demands that are placed on one another by meaning and understanding. Language is a means of communicating with others, by which we make our thoughts manifest to them. For an expression or a sentence to have meaning in this linguistic sense is for it to have shared significance among the members of a linguistic community in a way which facilitates the communicative exchange. To understand an expression or a sentence is to grasp the significance it has within the communicative exchange between the members of the linguistic community. McDowell speaks of this relation between meaning and understanding as being "contractual", in the sense that to understand an expression or a sentence is to come under an obligation to use it – "to judge and speak" – in certain determinate ways if one is to continue in the practice of using it with the same meaning.⁴ This introduces something that a would-be participant must latch on to in order to count as successfully integrated within the practice. McDowell writes, "If we regard an individual as aiming to speak a communal language, we take account of the possibility that he may go out of step with his fellows; thus we make room for an application of the notion of error, and so of right and wrong." (McDowell 1998f: 225) If a subject chooses to engage in the practice of a community, he thereby chooses to be

⁴ McDowell 1998f: 221

assessed by the norms that govern the practice. What counts as getting it right, then, is the ability to use the words in a way that accords with the dictates of the practice.

Understanding does not, on this picture, reduce to a subject's ability to interact with the members of a linguistic community, since someone might give the appearance of engaging in a communal linguistic practice even though they are in fact disposed to use the expressions of the community in a way that deviates in crucial respects from the members of the community. And even though these differences may elude detection in ordinary conversational circumstances, there would still be a clear sense in which such a person had not latched on to the correct practice of using those expressions, as defined by the norms of the community, and so failed to understand their meanings. I take it that this is one way of construing the intuition behind the famous Wittgensteinian concern about rule-following.⁵ Suppose we were to encounter a subject whose observable behaviour gave every indication that he understood the sense of 'add' that we denote by the '+' sign. Presented with a situation in which he has to add 12 and 7, he acknowledges the total to be 19; in a situation where he has to add 57 and 63, he offers 120 as the answer, and so on and so forth. To all intents and purposes, so far as any interaction with the subject is concerned, he appears to use the word 'add' correctly. But if we were to learn that once this subject encountered numbers greater than 10,000, which perhaps arise less frequently as numbers to be added in everyday contexts, he consistently used 'add' to denote a calculation on which 6 is deducted from the proper total (and it not simply that he *miscalculates* by getting his numbers mixed up; he quite deliberately goes through the process of subtracting 6)⁶ we would judge that 'add' as he uses it has a different meaning to that which it has for us, since we denote by 'add' a calculation which functions in a uniform manner irrespective of the values involved. To the extent that such a person sought to assimilate the practice of the linguistic community, to use expressions in the

⁵ Wittgenstein 1953 §185-201. For discussion of the concern and an interpretation that connects it to the idea of linguistic practice as I have used it here, see McDowell (1984).

⁶ In this respect, the subject's situation is supposed to mirror that envisaged by Wittgenstein, in which a subject learning to continue a series of +2 gets beyond 1000 and then starts counting 1000, 1004, 1008, 1012, and so on, not because of any computational failure on his part but because that is what he thinks the rule is supposed to be. Wittgenstein writes, "We say to him: "Look what you've done!" —He doesn't understand. We say: "You were meant to add two: look how you began the series!"—He answers: "Yes, isn't it right? I thought that was how I was *meant* to do it."—Or suppose he pointed to the series and said: But I went on in the same way." Wittgenstein 1953 §185

same way as other speakers and to be assessed by the norms that guide their practice, there is a clear sense in which we should say that that subject has failed to understand the meaning of the expression. In the relevant sense of understanding, the fact that the crucial deviancy eludes detection in ordinary conversational circumstances does not make it any less the case that the subject has not latched on to the correct meaning of the expression.

While it is certainly true that we who succeed in understanding one another may sometimes have discrepancies in the way we use certain words and yet, in spite of this, find ourselves able to enjoy some sort of successful interaction, we should perhaps resist the temptation to build our conception of understanding around such cases. For on this alternative picture, as we have already noted, to understand is to know the significance that an expression or a sentence has for the members of a linguistic practice. If we think of this in terms of language as a means by which speakers make manifest their thoughts to others, then understanding would seem to involve knowledge of which thought is expressed by a speaker's utterance. It is this sense of understanding we find in the following passages by Evans and Fricker:

Let us suppose that a speaker utters a sentence containing an expression which has a conventionally recognized information-invoking role, and it is clear that such a use is intended.... The audience must proceed beyond this, to the *right* (i.e. intended) interpretation. And if he is to be credited with understanding, he must *know* that it is the right interpretation. (Evans 1982: 310, italics in original)

When a human speaker-hearer of a language hears and understands an utterance by another speaker of a sentence of their shared language, she typically comes to know, or comes to be in a position to know, what speech act was effected in the utterance. (Fricker 2003: 325)

Understanding, in this sense, then, is intimately connected to knowledge. If this is right, then cases in which speakers interact with each other despite having crucial discrepancies in the meanings they attach to expressions cannot be the norm by which we illuminate our conception of understanding. To see this, consider a scenario where the circumstances of the deviant logicians or the deviant adder are much more widespread, so that it is no

longer just a couple of logicians with fairly academic concerns about one or two expressions of a language – rather, in this situation, large numbers of people who interact with each other have adopted idiosyncratic beliefs affecting their patterns of usage with respect to a considerable number of the words they use in their interactions. We might suppose that, as with the deviant logicians, these idiosyncrasies are far enough removed from the concerns of ordinary conversational discourse that they rarely – if ever – have cause to surface, so that interactions take place without the underlying discrepancies coming to the fore. There may be a sense in which we want to count what occurs between these speakers as understanding, since they are able to navigate their way through their everyday affairs without the differences in their meanings causing interference enough to be a stumbling block. However, to suppose that this sense of understanding could be exhaustive of those available would be to abandon the otherwise intuitive idea that understanding must be closely connected to knowledge of which thought is expressed by a speaker's utterance.

If a speaker gives expression to a thought, it is made available to be known by an audience who grasp the significance the speaker intends the words to have. An audience for whom the words have meanings crucially discrepant from that of the speaker cannot know (on the basis of the speaker's utterance) which thought the utterance was intended to express. The obvious obstacle lies in the audience's ability to get hold of the precise thought the speaker intended to communicate, since the words signify something different for the parties involved at either end of the communicative exchange. But even if it should turn out that the audience somehow latched on to the right thought, it seems counterintuitive to suppose that they have *knowledge* of which thought the speaker expressed. For, given that they assign different meanings to their words, the fact that they happen to get hold of the thought expressed by the speaker would seem be little more than a fortuitous accident of the sort we ordinarily take to be incompatible with knowledge. Heck, writing about a different problem, seems to make a similar point when he says,

If the content one speaker forms in reaction to an assertion bears no particular relationship to the content of the belief the latter was expressing, then, even if her newly formed belief were true, she should merely have been *lucky*. Again, if the

beliefs we form in reaction to sentences asserted by others bears no close relationship to the meanings of those sentences, then, even if those beliefs are true, we would not, in ordinary cases, *know* them to be true. We would not even have acquired justified beliefs. (Heck 1995: 91, italics in original)

Heck's remark is made in the context of discussing an audience's capacity to know *that* the content expressed by a speaker is true. But it seems to have equal application to an audience's ability to know *which* content or thought is expressed by a speaker. If a subject is to know that some specific thought is the one to which a speaker has given expression, then it cannot be by mere accident that the subject gets hold of the right thought. Given this, if we are to preserve the connection between understanding and knowledge, then cases where there are significant discrepancies in the ways speakers use words cannot be the norm by which we illuminate our conception of understanding.

The issue here is whether there is something we presuppose a subject to have achieved – something we expect the subject to have latched on to – when we attribute understanding of an expression or a sentence. Williamson, we have seen, argues that there is not and that understanding is therefore unable to perform the epistemic role necessary to support an account of analyticity. However, if we are to make room for the idea of language as a genuinely public phenomenon that enables the members of a linguistic community to share in the exchange of thoughts by exploiting the mutual significance of expressions and sentences, then we should be dissatisfied with the suggestion that the picture presented by Williamson exhausts what it is to understand a language. For part of what it is to understand, on this alternative picture, is to grasp the significance that expressions and sentences of the language have for members of the practice. For example, a speaker who uses an indicative sentence gives expression to a thought about the world and its being a certain way. It seems that we should count a subject as having understood the speaker only if the subject knows (on the basis of the speaker's utterance) which thought the speaker expressed, that is, if he knows how the speaker represented the world to be.

(iv) High-level theoretical concerns and disagreement

Does the foregoing preclude counting Williamson's deviant logicians as understanding 'All vixens are female foxes'? If it does, the worry would be that all we have achieved is to redefine what it is to understand in a way that reinstates the stronger constitutive link between understanding an analytic truth and knowing it in order to circumvent Williamson's concern. If, on the other hand, the deviant logicians count as understanding in the sense outlined in the previous section, then we would seem to have identified a sense of understanding stronger than that presented by Williamson – and therefore at least potentially capable of shouldering the epistemic work required for analyticity – while accommodating his insistence that it is always possible for a subject to understand a putatively analytic truth and yet fail to know it.

It is helpful to consider the difference between those who deviate from the norm of the linguistic community in their use of an expression in cases like that of the deviant adder, discussed in the previous section, and those who do so in cases like those of Peter and Stephen. A case such as that of the adder looks like a clear case of a failure to understand. There is a pattern of behaviour to which the adder regards himself as trying to conform; he attempts to assimilate the norm that guides the practice of others in the linguistic community. However, his behaviour in attempting to add values greater than 10,000 betrays the fact that he has failed in the effort to be guided in the way that other users of 'add' are. There is a mismatch between what he supposes the practice of the community to be and the practice as it really is.

It is not clear that there is a mismatch of this kind in the cases of the deviant logicians. As Williamson describes them, Peter and Stephen may be professors of philosophy who have defended their peculiar logical views in leading refereed journals.⁷ The point of this background story is that there is nothing one could say to them to simply help clarify or correct our practice of using the relevant logical expressions. Whereas it may illuminate the deviant adder to point out that his use of 'add' does not accord with the norm, it would not, for example, be illuminating to Peter to be told that we do not use the universal quantifier in such a way as for it to be existentially committing. In having

⁷ Williamson 2006: 12 and 2007: 89

defended their views, they must already appreciate the way in which they diverge from common practice. That is, they know what the common practice is in relation to the relevant expressions but have some reason for adopting a difference pattern of use. This is crucial in assessing their understanding. On the stronger conception of understanding outlined above, it is a condition of counting an audience as having understood a speaker that the audience knows which thought the speaker has intended to communicate – how, in the case of an indicative sentence, the world has been represented to be by the utterance. Peter and Stephen certainly seem to satisfy this condition. For their ability to defend the unorthodoxy of their views presupposes knowledge of what the prevailing orthodoxy is. So, when a speaker who does not share their pet theories gives expression to a thought using the relevant logical expressions, Peter and Stephen are nevertheless able to know what the speaker sought to communicate by the utterance and thus live up their part in the communicative exchange. There is therefore no mismatch between what they suppose the practice of the linguistic community to be and the practice as it really is, since they know how other people use the relevant logical expressions, although they believe there to be a defect in the prevalent practice that motivates using the expressions as they do. (I say more about what we should understand them to be advocating in relation to the practice in the next chapter.) This suggests that they should count as understanding the sentence ‘All vixens are female foxes’.

One might worry that there must be a sense in which Peter and Stephen fail to understand on the grounds that a thought is the kind of thing that has objective truth-conditions yet, in the case of the thought expressed by ‘All vixens are female foxes’, Peter and Stephen fail to know them (since they assign different truth-conditions to the sentence).⁸ But this does not follow. Suppose I have a friend who I know to be prone to getting confused about the difference between the meanings of ‘immoral’ and ‘amoral’, so that he sometimes uses one when he clearly means the other. If, in the context of discussing the impermissibility of uneven wealth distribution, my friend were to assert ‘Capitalism is amoral’, I may take him to have fallen prey to his confusion and interpret him instead to have been expressing a different thought, namely, that capitalism is immoral. My prior knowledge of his susceptibility to confusion (together, of course, with

⁸ A version of this line of argument was put to me by Marcus Giaquinto.

help from the conversational context) means that despite my friend's misuse of the word, I am nevertheless able to know which thought he intended to express by his utterance. So it is for Peter and Stephen. In their eyes, there is something defective about our practice of using the relevant logical expressions, so that the meaning we attach to the sentence 'All vixens are female foxes' is in some way incorrect. Although the truth-conditions they assign to the sentence are different to those we take ourselves to express, their prior knowledge of our "error" means that they are nevertheless able to know which thought we intend to express by our use of the sentence.

This situation where a subject has full knowledge of what the ordinary practice is but has qualms about the correctness of that practice would seem to characterise what Williamson in one of the quoted passages above refers to as "high-level theoretical concerns." They are theoretical insofar as they are sufficiently far removed from everyday concerns as not to interfere with the practicality of people like Peter and Stephen being able to engage with other speakers, high-level insofar as they are founded upon a prior knowledge of the existing practice together with an appreciation of how and why they deviate from this. Thus, to have a high-level theoretical concern entails that a subject has an understanding of that about which he has a concern.

In his most recent writing on the matter, Williamson goes further than this, suggesting that what is true of the deviant logicians may be extended to ordinary speakers who dissent from a truth such as 'All vixens are female foxes' without possessing the kind of theoretical concerns exhibited by Peter and Stephen. He writes,

We can also imagine untheoretical native speakers whose unreflective patterns of assent and dissent to non-metalinguistic sentences are those which Peter [and] Stephen...respectively recommend, although they lack the reflective capacity to rationalize those patterns by appeal to formal semantic theories. They too would be able to fit in well enough with the rest of the linguistic community, to engage smoothly in useful communication and adjust to their differences with other speakers in order not to attract too much attention. (...) Once we concede that Peter [and] Stephen...are competent speakers, we can hardly refuse the same classification to

other speakers merely on grounds on their unacquaintance with formal semantics.
(Williamson 2007: 99)

Here I think we may draw a line and simply deny Williamson his claim. For the question that becomes pressing for this view is: what is the difference between those who, on the one hand, dissent from such a truth and still count as understanding it and, on the other, those who dissent through a failure to have understood one or more of the constituent expressions? Williamson's answer to this, as can be seen from the above passage, is that the difference consists in the ability of those who understand to successfully interact with the rest of the linguistic community. However, once we acknowledge that a subject might enjoy smooth interaction with other speakers while failing to count as understanding in a fuller sense, the question regains its bite.

If we accept the argument of the previous sections, the difference between those who dissent from something like 'All vixens are female foxes' despite understanding it and those who dissent precisely because they fail to understand it lies in the ability of the former group to know which thought non-dissenting, ordinary speakers take themselves to express by their use of the sentence. In these instances, a subject should know what is the common practice of using the expressions and in what respect they deviate from that practice. This is not, as Williamson insinuates, to demand that anyone who chooses to dissent must be acquainted with formal semantics. It is, however, to insist that dissent is incompatible with understanding if it stems from a simple failure to grasp the communal practice of using the expressions. Once understanding in this sense is achieved, a subject's dissent will then reflect something approaching a "high-level theoretical concern," even though the concern itself need not be especially theoretical or born out of any preoccupation with the finer points of formal semantics – it may be theoretical only insofar as it does not interfere with the practicality of communication. But it must accompany a recognition both of what the accepted practice is and in what way the subject's own usage deviates from that norm. Since we have reason to suppose that Peter's and Stephen's dissent is of this kind, we may count them as understanding 'All vixens are female foxes', even on this stronger picture of understanding.

(v) Concluding remarks

Williamson argues that the appeal to *being in a position to know* is unable to preserve an account of analyticity from the difficulties introduced by cases of disagreement. Once we recognise the possibility that for any candidate analytic truth someone could understand it and fail to know it, he argues, we are restricted to a picture of understanding too weak to do the epistemic work required for analyticity. He supposes that we must view understanding as requiring only the ability to get by in our interactions with other speakers; consequently, there is nothing a subject is required to accept in order to count as understanding any given utterance.

We can now see the problem in Williamson's argument. The picture of understanding he thinks we are forced to accept in order to accommodate cases of disagreement is in fact too weak to do justice to the idea that understanding involves knowing which thought is expressed by a speaker's utterance. In order to make room for this, we need a conception of understanding stronger than that presented by Williamson. Yet on the stronger account sketched here, we retain the ability to accommodate the possibility that for any given truth, someone may understand it and nevertheless harbour reasons to dissent from it. So the possibility of such disagreement does not require us to accept the weaker conception of understanding presented by Williamson. In the next chapter I offer an account of how the stronger picture of understanding sketched here enables a line of explanation for how understanding could explain our knowledge of analytic truths.

How Could Understanding Explain Our Knowledge of Analytic Truths?

We have reached a stage where some positive account in answer to the above question is required. The question is an instance of a more general kind labelled by Cassam as *how-possible* questions. A how-possible question is one that seeks to address how something which has the appearance of being *impossible*, given certain other things one knows or believes, is nevertheless possible. Cassam observes that many philosophical problems take this form: we start from the assumption that something is the case but then encounter some real or apparent obstacle that creates a puzzle for how things could be as we supposed.¹ In raising the corresponding how-possible question, we aim to resolve the puzzle, perhaps by showing that there is no real conflict or by demonstrating that the obstacle can be overcome.

An epistemological how-possible question is one that takes as its starting point some epistemic achievement which we suppose to be genuine and asks how that achievement is possible in the face of some reason to doubt that an achievement of this sort could be realised. For example, we might worry about the possibility of acquiring knowledge by testimony if we suppose that the acquisition of knowledge by this method requires a subject to have ready evidence concerning the reliability of the person from whom it was acquired. The achievement from which the how-possible question would proceed in this case would be the knowledge we acquire through testimony on numerous occasions on a day-to-day basis. The obstacle is the worry that in many cases we seem to accept testimony without having much to go on regarding the reliability of our informant. In addressing the how-possible question, we seek to show how, despite the apparent obstacle, we are able to have knowledge of the kind in question.²

Our present interest is in knowledge acquired on the basis of understanding. The achievement from which we start is the seeming difference in the way we know, for

¹ Cassam 2007: 3-4

² A different sort of answer, of course, would be one that showed the obstacle to be real and insurmountable, with the consequence that despite initial appearances we do not in fact have knowledge of the kind in question.

example, the proposition that all vixens are female foxes and the proposition that vixens have dichromatic vision. Whereas the latter is fairly typical of the knowledge we acquire insofar as it results from investigation into how things are in the world, the former does not have the appearance of being known in this way. Rather, we are commonly inclined to suppose that our knowledge in these cases is somehow to be explained by our understanding of the expressions involved or by our grasp of the relevant concepts. The identification of our understanding as the means by which we know truths of this general kind is the first stage of the how-possible response. The second stage consists in resolving the puzzle created by the reasons for thinking that the achievement is illusory. We have seen in the preceding chapters that there have been various obstacles perceived to impede the idea of knowledge acquired by understanding, from Quine's objections against Carnap to Williamson's multiple concerns centred around the theme of the need to accommodate disagreement about truths of this kind. I have argued that the defender of analyticity has the resources to respond to these obstacles, leaving intact the claim that there are truths our understanding of which places us in a position to know them.

The tripartite structure of a how-possible answer is completed by a third stage, aimed at providing a more positive account of how it is that the means identified at the first stage could be a source of knowledge. The task is to identify what Cassam calls the "enabling conditions" for the relevant kind of knowledge, the conditions under which it is possible to acquire knowledge by the particular method in question. The demand for this third stage reflects the thought that "there is more to explaining how something is possible than showing that it isn't impossible." (Cassam 2007: 9) A positive account along these lines is perhaps all the more important in the case of an unobvious source of knowledge like that of understanding. For there are some human faculties that just look like good candidates for being knowledge-conducive, such as that of visual perception, whose very role is to furnish us with information about the world around us. Understanding, in general, is not like this. After all, there are many truths I understand: I understand the sentences 'Snow is white' and 'Cambridge is north of London', but in neither instance do we have the inclination to cite my understanding as the primary explanation for my knowledge of the propositions they respectively express. Even in cases of testimony, where arguably understanding has a greater explanatory role, we might still be reluctant to think that the

knowledge I acquire is explained purely in terms of my understanding the words of an informer. So understanding has various roles to play in our knowledge acquisition, at least some of which do not involve it being the primary source by which we would explain the knowledge we acquire. It is therefore incumbent upon a satisfactory defence of analyticity to provide something more by way of a positive explanation of how understanding could play the epistemic role it requires.

The defender of analyticity can go about such a task in broadly two ways. On one hand, we might seek to divide the general class of analytic truths into further sub-divisions, providing an explanation specific to each as to how understanding could explain a subject's knowledge of the truths belonging to these.³ On the other hand, we might seek a more general explanation of the role that understanding plays in knowing analytic truths *per se*.⁴ The two approaches are complementary rather than in opposition with each other. We might start out looking for a general account of how understanding could perform the epistemic role required for analyticity and then show how this general account is realised for different kinds of analytic truths. Alternatively, we might begin with the particular instances and then abstract away from these to isolate a more general picture about the role that understanding plays. Here I shall focus on outlining an explanation of the more general kind. The scale of this task, Peacocke observes, is greater or lesser in accordance with the extravagance of the conception of understanding from which we start out. If we have a view like that which Peacocke takes to be commonly attributed to Gödel, on which understanding offers some sort of direct insight or contact with meanings or concepts,⁵ so that on occasion one could be said to somehow 'see' that an analytic statement is true, the explanatory task is made a little easier (though at the risk of appearing a little less plausible too).⁶ Peacocke writes,

The more sober the view of what is involved in attaining a priori knowledge, the more challenging the task of explaining how understanding has the epistemological

³ For one example of the sub-divisions one might draw, see Boghossian 1997: 334. An exemplary statement of how to apply this method is to be found in Peacocke's discussion of conceptual truths involving colour properties in his (2000).

⁴ Peacocke (1993) is perhaps one of the better known efforts that fall into this category.

⁵ Peacocke 2000: 261-2

⁶ Strictly speaking, of course, an account that veers too far in this direction becomes less recognisably an account of analyticity and more, perhaps, one of rational intuition.

power. The more sober rationalist has to account for the epistemological phenomena from a non-exotic theory of understanding or concept-possession. (Peacocke 2000: 257)

The truth of this remark is supported by Williamson's account of understanding which was discussed in the previous chapter. Williamson's conception of understanding is decidedly non-exotic and unextravagant, and if it were to exhaust the conditions for understanding then it would be very difficult indeed to see how understanding should play the epistemic role required for analyticity. What we need, then, is a satisfactory general conception of understanding, together with an account of how this would enable it to explain our knowledge of analytic truths. The simplest, and perhaps most attractive, claim for the defender of analyticity to make is that the conditions for understanding just are those that suffice to put a subject in a position to know an analytic truth. The difficulty then becomes one of showing that this provides sufficient scope for understanding to have the necessary explanatory power.

In the previous chapter I argued that we need not accept Williamson's picture on which all there is to understanding is the ability to interact with other speakers of a language, arguing that if we are to preserve the connection between understanding and knowledge of what is said, we must have a stronger sense of understanding according to which a subject is able to know precisely which thought a speaker intends to express by an utterance. If a speaker asserts a claim using an indicative sentence, thereby representing the world to be some way, a subject who counts as understanding in the relevant sense will know how the speaker represents the world to be, that is, under what conditions the proposition asserted by the speaker is true. In this chapter I pursue this idea further. In particular, I argue that it is part of being a competent language user that a subject will have personal-level knowledge of truths-conditions that is made available through the phenomenology of understanding an utterance. In thus knowing how the world is represented to be, a subject must have some grasp of the possibilities that are ruled in or ruled out by the truth of the relevant proposition. We can then apply this general picture of understanding to explain our knowledge of analytic truths. For an analytic truth, being knowable independently of experience, must be one whose constituent expressions in

some sense guarantee the fulfilment of its truth-conditions. Thus, to understand an analytic truth is to be in a position to recognise that its truth-conditions must be fulfilled, thereby entitling a subject to endorse it.

There are two principal claims here in need of further examination:

(1) Understanding an utterance involves knowledge of its truth-conditions.

and

(2) Knowing the truth-conditions of an utterance involves having some grasp of the possibilities that are included or excluded by the truth of the proposition expressed.

If (1) and (2) can be defended, then assuming (3),

(3) An analytic truth is one whose truth-conditions are guaranteed to be fulfilled,

then we can make a case for (4),

(4) Understanding an analytic truth enables a subject to recognise that its truth-conditions must be fulfilled.

In the first two sections of this chapter I defend the principal claims (1) and (2). In section (i) I consider the claim that understanding an utterance involves knowledge of its truth-conditions. Although the view has had its detractors in recent literature, I argue that the arguments presented fail to tell against it, at least on the reading of the claim relevant to present purposes. In section (ii) I argue that a subject who knows the truth-conditions of an utterance must have some grasp of the possibilities that are included or excluded by the truth of the proposition. With these two claims, we seem to have a broad outline of how understanding could explain our knowledge of analytic truths. However, there may still be a worry that the account outlined is no longer able to accommodate the kind of defeasibility that made the idea of *being in a position to know* so central to the previous

chapters. The concern might be that if understanding involves knowledge of truth-conditions, so that understanding an analytic truth enables one to recognise its truth-conditions to be fulfilled, then what room is there for the possibility that a subject might understand it and yet fail to know it? It may appear that we have built into the idea of understanding an analytic truth that one must thereby also know it, regressing to the constitutive model of analyticity from which we were trying to get away. But this conclusion is not one we are forced to accept. In section (iii) I show that the account outlined in the preceding sections has much in common with the so-called ‘metaphysical’ conception of analyticity advocated by the likes of Carnap, with the practice of a linguistic community forming a framework very similar to that which we saw to be so central to the Wittgensteinian picture of the positivists. In section (iv), using this model of the framework, I demonstrate that there is room for the view defended here to accommodate disagreement about analytic truths.

(i) Understanding and knowledge of truth-conditions

There is a neat and attractive argument we can sketch to motivate the claim that understanding involves knowledge of truth-conditions. A speaker who uses an indicative sentence expresses a thought about the world and its being a certain way. The meaning of the sentence delimits a way for the world to be, thus imposing a condition on how things must be in the world if the utterance is to be true. Understanding the speaker requires knowing what thought is expressed by their utterance. A subject who hears the utterance counts as grasping the thought expressed by the speaker if he thereby knows how the world has been represented to be. A subject who hears the utterance but fails to know how the world is represented to be would not count as having grasped the thought expressed by the speaker. Understanding, therefore, involves knowing how the world is represented to be by the speaker’s utterance. Since how the world is represented to be is the condition by which the utterance of the sentence is assessed for truth or falsity, understanding involves knowing the truth-condition of the utterance of the sentence.

Matters may not be quite so straightforward, though it certainly offers an attractive enough starting point. However, the view that understanding involves knowledge of truth-conditions is one that has come under challenge in some of the literature on

understanding. One reason for this, as we shall see, is that the slogan “understanding involves knowledge of truth-conditions” has connotations of an alternative view, different from that for which I am arguing, and there is evidence of a degree of confusion about the respective commitments of these different views. But there are also less misguided objections that need to be addressed. I shall examine two lines of argument that threaten the claim, one articulated by Soames (1989, 1992) and another put forward by both Hunter (1997) and Fricker (2003). Soames, we will see, does not take himself to be arguing against the claim I am making, but the argument he addresses is instructive for bringing out the confusions that arise from the different commitments associated with the slogan.

Soames (1989, 1992)

Our starting point should be the distinction that theories of understanding aim to capture, viz. the difference between those who understand a language and those who do not. This is, at least in part, a phenomenological issue. When a speaker utters a sentence of their language, they emit a sound. To those who do not share the language of the speaker, they are – in some sense – just sounds. A hearer might not share the language yet know the sounds to be meaningful, perhaps because the responses of others around them indicate that the sounds are meaningful to them. But phenomenologically, in such a scenario, the utterance does not directly strike the hearer as meaningful; it remains, for such a hearer, mere noise. This serves to remind us, says Smith, that “at one level all that goes on in the environment when we speak is the issuing of certain sounds.” (Smith 2006: 949) For those who share the language, however, the sounds will not register as mere noise. Rather, they are experienced as a contentful speech act, imbued with meaning. There is not first a hearing of some sound and then some separate act of grasping its semantic significance, but instead the experience of being confronted with discrete words in an ordered, grammatical form, indistinguishable from the experience of hearing those sounds that remain available to the understanding and non-understanding alike. Fricker writes,

[One who understands an utterance] hears the utterance not merely as sound, but as the speech act that it is. The representation of meaning is not phenomenologically distinct from her hearing of the sounds. She hears its meaning in the utterance itself, experiences it as a semantically laden event. It is for her as if, in perceiving the

utterance, she perceived not just the sounds, but equally perceives their meaning. It is a fact of phenomenology that we enjoy such understanding-experiences, quasi-perceptions of meaning (Fricker 2003: 325)

It seems reasonable to suppose that the phenomenological difference here is explained by there being something that is known by a fellow speaker of the language that is not known by someone who does not share the language. The most obvious candidate for what is known is the meaning of the uttered sentence: when a speaker of English utters the sentence 'Snow is white', a hearer who shares the language and understands the speaker will know not just which words the speaker used but also what those words – arranged in their particular order – mean, specifically, in this instance, that snow is white. But there is a stronger claim that has been pervasive in the literature to the effect that the competence exhibited by one who understands is knowledge of the theorems delivered by a semantic theory. The idea is that we are capable of constructing theories of meaning that offer a description of the use to which speakers of a language put their words. They may describe, for example, the meanings of individual expressions, how these can be combined with one another, and what they may be used to say when combined in certain ways. Since these theorems in some way describe the practice of the speakers of language, we might suppose that there must be some relation by which we can connect the pronouncements of the theory with what goes on inside the heads of the speakers. As Smith puts it,

If these theoretical descriptions are not to remain blankly external to the mind of the speaker they must have some connection with the concept of knowledge. For either the theory is envisaged as describing, or reflecting, knowledge speakers actually have, consciously or not, or it concerns a body of knowledge that would enable someone in possession of it to use and understand the language. (Smith 1992: 112)

So the thought is that a semantic theory not only offers us a description of speakers' practice but serves doubly as a theory of competence, that is, a theory of what is known by those who understand the language. And since the time Davidson first noted that the representational properties of semantic information could be captured by using Tarski-style T-sentences, the notion of truth-conditions has been a popular one amongst

semantic theorists. That the meaning of a sentence – the role it has in the linguistic practice of speakers in a language – might be given by a T-theorem of the form $\langle S \text{ is true iff } p \rangle$ is accompanied by the supposition that knowledge of the T-theorem is what underlies the competence exhibited by those who understand the language.

On one view that emerges from this picture, the claim that understanding involves knowledge of truth-conditions is the strong claim that one understands a language only if one counts in some way as knowing the T-theorems by which we represent the meanings of sentences in the language. This introduces a considerably greater commitment than seemed to be involved when we initially sketched the motivation for supposing that understanding must involve knowledge of truth-conditions. For the T-theorems in a semantic theory of the envisaged sort employ a fairly technical apparatus, drawing on concepts that many competent, ordinary speakers may not recognise to be within their grasp, even once it was pointed out to them that they reflected an aspect of their own practice. This is the basis of Soames' objection. For the T-theorems of a semantic theory give an irreducible role to the concept of *truth*, and so the claim that a subject has knowledge of truth-conditions (in something like their T-theorem form) entails that the subject be in possession of the relevant concepts to formulate the T-theorem, including that of *truth*. Soames says,

Knowledge of truth-conditions...presupposes possession of a metalinguistic concept of truth. Thus, the claim that such knowledge is necessary for understanding meaning entails that no one can learn or understand a language without first having such a concept. But this consequence seems false. Certainly, young children and unsophisticated adults can understand lots of sentence without understanding 'true', or any corresponding predicate. (Soames 1989: 578)

If this is right, the strong reading of the claim that understanding involves knowledge of truth-conditions is too strong: it is no part of our ordinary notion of understanding that a subject needs to possess the concepts required by the apparatus of a semantic theory. But this strong reading was always committed to more than we needed to get from the basic idea that understanding involves knowledge of truth-conditions. All we need is that, in

understanding the utterance of an indicative sentence, a subject thereby grasps how the world is represented to be. So, in understanding an utterance of 'Snow is white', a subject should grasp that the world is represented to be such that snow is white. We can call this 'knowledge of truth-conditions' because what is known by the subject is some way for the world to be, as delimited by the sentence, which is, necessarily, how the world is if the proposition expressed is true. This imposes no requirement that the subject himself should think of what he knows in these terms. This weaker sense in which a subject may have knowledge of truth-conditions not only evades Soames' concern but seems also to be one to which Soames himself is broadly sympathetic; as he writes at one point,

The child will get along fine so long as he knows that 'Momma is working' is to be assertively uttered only if Momma is working; 'Daddy is asleep' is to be assertively uttered only if Daddy is asleep; and so on. The child doesn't have to say or think to himself, 'There is a general (but defeasible) expectation that for all x, if x is a sentence, then one is to assertively utter x only if x is true.' (Soames 1989: 578-9)

Now, Soames' characterisation of his opponent's stance here is something of a caricature. There is little reason to suppose that someone like Davidson (Soames' primary target) is committed to the idea that any subject, child or adult, should think about things in this abstruse way.⁷ Nonetheless, we do face an important question concerning what to say about the access a subject has to the rules or theorems of the theory by which we might characterise the meanings of the sentences in his language. It was, after all, the idea articulated in the quotation from Smith that the pronouncements of a semantic theory must be in some way known by speakers which led us to the strong version of the claim that understanding involves knowledge of truth-conditions. If we are to reject the strong reading of the claim as too strong, it is desirable that we should have some story to tell about the relation between semantic theorems and what is known by speakers in a way that does not mandate acceptance of the strong claim. What we need, then, is some

⁷ Precisely what Soames himself thinks about this is not so clear. He says, "It is one thing to suppose that all speakers of English know that 'Snow is white' is true in English iff snow is white.... It is quite another to maintain that they must know that which is stated by the Tarski-style theoretical apparatus.... But this theoretical knowledge is precisely the sort that they have to have, if knowledge of that which is stated by the usual sorts of translational truth theories is to be necessary for understanding a language." (Soames 1992: 22) However, this isn't the claim to which Davidson is primarily committed, since Davidson claimed that knowledge of the semantic theory would be sufficient for someone to understand a language. Cf. Davidson (2001).

account of how we can respect the demand that there be some connection between our semantic theorems and what is known to speakers, without requiring that speakers actually possess the concepts deployed in constructing a semantic theory.

We typically think of a subject's knowing something as being what we might describe as ordinary, propositional knowledge. It is the sort of knowledge we would often expect a subject to know that he has and, as such, comprises concepts grasped by the subject. However, if we can be said to in some sense *know* the theorems of a semantic theory, it seems the knowledge in question cannot be of this ordinary sort. It is not the kind of thing of which ordinary speakers are typically aware, and might only be able to be articulated by drawing on concepts not grasped by the subject himself. So if we wish to call what the subject has here *knowledge*, in order to reflect its status as that which explains the difference between the speaker who understands and the non-speaker who does not, then we may wish to think of it as knowledge of a different kind. We find this idea in the work of Chomsky, who talks about our knowledge of language as being *tacit* knowledge, and in Dummett, who prefers to speak of *implicit* knowledge (which is not to say that these amount to the same thing in the eyes of their respective advocates, nor that either lacks points of contention).⁸

We can further motivate thinking of our knowledge of language as a kind of tacit knowledge by reflecting on the phenomenological considerations that we took as our starting point for understanding. We said that when a subject understands an utterance, it is experienced as something inherently contentful, that there is something like a direct access to the meaning of the uttered sentence. Now, we know that strings of words are not themselves inherently contentful; there are plenty of combinations of words I might put together that would be quite meaningless to any fellow speaker of English. So, rather, a string of words has meaning for a subject by virtue of its having some structure which transforms it from a mere concatenation of words to something capable of being experienced by an audience as – to use Fricker's phrase – a semantically laden event. We might be able to uncover the structure of a sentence by articulating rules about things like the permissible modes of combination, etc., but the structure itself is no part of the

⁸ See Chomsky (1986) and Dummett (1993)

phenomenological experience of understanding: our direct access is to the content itself, not the structure by virtue of which there is content. Insofar as we are interested in the personal-level knowledge that is made available to a subject by an understanding experience, we may draw a distinction between that of which a subject is aware in understanding an utterance and the capacities that contribute to the subject having that awareness.

The idea I am gesturing towards is one written about by McDowell in discussing the personal/sub-personal distinction introduced by Dennett. McDowell writes,

The idea is that “events within us” are contentful *by virtue of* their structure; for the content possessed by an internal event or state is a function of its function in the organism, and this function “is – in the end, must be – a function of the *structure* of the state or event and systems of which it is a part.” But only some of the events and states within us that possess content in this way make available to *us* the content that their structurally determined role in the system confers on them (although they transmit freely among themselves, unnoticed by us). And when they make their content available to us, the structure in virtue of which they have it does not figure in an accurate phenomenology of our consciousness; it remains a topic for theory, not mere introspective noting. (McDowell 1998g: 341-2, italics in original)⁹

Content here is that of which a subject has direct awareness at the personal level, and draws upon concepts grasped by the subject. *Structure* is that in virtue of which content is made available to the subject, though is not itself made directly available to the subject in conscious experience. Because of this, we may think of it as being at a sub-personal level, capable of being described using concepts not grasped by the subject. Applied to our knowledge of language (for that is not McDowell’s concern in the above passage), that of which a subject is directly aware in understanding an utterance is the meaning of the sentence, which represents the world to be a certain way. The structure is that in virtue of which the utterance of the sentence comes to have that contentful significance for the subject – how it gets to represent the world as being some particular way. What we get

⁹ The quotation marks refer to remarks by Dennett in his paper ‘Toward a Theory of Cognitive Consciousness’.

from a semantic theory are claims about the structure of sentences, what they are used to say, and how they represent the world. That is, it seeks to explain the meaningfulness that utterances have for us at the contentful, personal level. In this way, we may view a semantic theory as aiming to describe, or represent, aspects of what goes on at the sub-personal level of a subject's language-processing system to give rise to the personal-level experience of understanding an utterance. This division offers a means by which we might connect the pronouncements of a semantic theory with what goes on inside the heads of speakers without requiring that speakers themselves grasp the concepts employed in constructing the theory. If we want to call what a subject has in relation to these semantic theorems *knowledge*, then the lack of direct awareness and the conceptual sophistication required to describe it certainly motivates thinking of it as a distinct kind of knowledge.¹⁰

The upshot of all this is it is possible to take Soames' point that understanding does not require knowledge of truth-conditions in the strong reading of this claim on which we have ordinary, personal-level knowledge of the T-theorems laid down in a semantic theory. We can grant that we are cognitively related to these theorems, so they are not "blankly external" to the minds of speakers; we may also grant that they have a role in explaining part of what goes on in the process of understanding. But this does not commit us to the claim that a subject need have any awareness of the theorems, or even possess the concepts necessary to formulate them. The distinction between personal-level content and sub-personal structure creates room for the idea that, at the personal level, understanding an utterance of a sentence involves knowledge of truth-conditions in the weaker sense of having an awareness of how the world is represented to be.

Hunter (1997), Fricker (2003)

Although Soames' objection to the claim that understanding involves knowledge of truth-conditions targets a different conception of what this amounts to from that which we initially sought to motivate, the argument is nevertheless of interest at least insofar as it

¹⁰ There are various issues surrounding whether we really ought to think of the relation as one of knowledge, with everything that entails from a philosophical standpoint, or whether we should think of it in less normative terms. I will not attempt to adjudicate on these issues here. The important point for present purposes is that we can make room for the possibility that we are somehow related to the theorems of a semantic theory, however we want to characterise that relation.

helps us to draw out the divergent commitments associated with the view. However, even when extricated from the stronger commitments concerning subjects' knowledge of T-theorems, the claim that understanding involves knowledge of truth-conditions is not removed from contention. Both Hunter and Fricker argue that the phenomenological character of understanding is akin to that of visual perception. Since, they argue, states of perception are distinct from any state of belief or knowledge, so too are states of understanding. Consequently, understanding, in and of itself, cannot involve knowledge of truth-conditions.

Hunter and Fricker are interested in how the experience of understanding an utterance could ground a subject's knowledge of what is expressed by a speaker. On the view I have been advocating, knowing what is expressed by a speaker is simply a part of what it is to count as understanding an utterance. Hunter makes clear his dissatisfaction with this answer in the following passage:

To assume that states of understanding are states of knowledge or belief, that to understand a text or speech act is to have already incurred a commitment concerning its meaning, is to prohibit understanding from playing a role in justifying such commitments. Far from solving the epistemological problem of how meaning is known, this assumption simply raises it anew as the question [of] how the commitments that constitute understanding are justified. Indeed, it seems to exacerbate the problem since it is not at all clear where one might look for such a justification. (Hunter 1998: 577-8)

So, the thought goes, when a speaker produces an utterance, we form a belief about what was said by the speaker, which in typical cases we suppose to be an instance of knowledge. In order to achieve this status, those beliefs must be in some way justified. The obvious candidate to cite as the source of this justification is the experience of understanding the speaker; as Hunter intimates, it's not clear where else we might look. Hunter's worry is that if understanding itself involves knowledge of what is expressed by a speaker, then it is no longer capable of having the justificatory role he envisages, since there would need to be some further source of justification for the knowledge inherent in understanding. Hunter supposes that since we *do* have knowledge of the things expressed by speakers,

understanding must play the justificatory role, which in turn means that it cannot – in and of itself – involve knowledge of what is expressed. We find a very similar picture of the relation between understanding and knowledge of what is said in Fricker, who says,

[The] quasi-perception of the heard utterance as a meaningful speech act...provides the hearer with a ground, a justifying basis, for her belief about what has been said, which renders it knowledgeable.... (Fricker 2003: 341)

Perhaps somewhat ironically, Hunter and Fricker both motivate this view by invoking the kind of perceptual defeasibility we examined in the earlier chapters in discussing the idea of *being in a position to know*. The scenario described was one in which there is a machine capable of producing perfect hallucinations of an orange together with a subject who is aware of the machine's capability, knowing that he would have exactly the same phenomenological experience whether he were to be confronted with an orange or the mere appearance as of an orange in front of him. Unbeknown to the subject, the machine is broken, so that if he is having the experience of seeing an orange, it can only be because there is an orange in front of him. When confronted with the orange, he has a veridical experience with the content that there is an orange in front of him. Yet the subject's knowledge of machine's capabilities cause him to refrain from endorsing that content. In such a scenario we have a genuine state of perception, providing information about things in the world to a subject, without any accompanying belief on the part of the subject that the world is as it is represented to be. What this shows, Hunter argues, is that "states of perception belong in a fundamentally different epistemic category from states of belief and knowledge:" (Hunter 1997: 565)

states of perception...are not states of belief. Rather, perception ...[is] a source or store of information on which a subject might base a belief. Subjects are not compelled to believe the evidence supplied by perception.... (Hunter 1998: 572)

Hunter and Fricker point out that it is possible to generate analogous scenarios for states of understanding. Suppose two people are standing in a crowded and very noisy area, trying to talk with one another. One might shout through the noise, say, that Cameron is

Prime Minister. The other might hear the assertion properly and understand his friend to have said that Cameron is Prime Minister yet, suspecting that the jostling and the noise around them is causing auditory interference, decide to withhold judgement about whether that is in fact what was said.¹¹ So as with the perceptual case, we have a genuine case of understanding, providing information to a subject regarding what was expressed by an utterance, without any accompanying belief on the part of the subject that what he understood to have been expressed was in fact what was expressed by the utterance. As a result, Hunter and Fricker claim, we should apply the lesson of the perceptual case to that of understanding:

states of understanding are like states of perception in being states of awareness that can serve as evidence on which to ground beliefs (...) it is possible for a subject not to believe that a piece of text or speech act mean what she understands it to mean. (Hunter 1998: 570)

states of understanding, like visual experiences, exhibit belief-independence. That is to say, a hearer's quasi-perception of content and force may persist, despite her lack of belief that what she apprehends is indeed being said. (Fricker 2003: 345)

If understanding is entirely distinct from belief or knowledge in this way, then our knowledge of what is expressed by utterances must be distinct from our states of understanding. Understanding, then, would not involve knowledge of truth-conditions. Note that the relevant construal of *knowledge of truth-conditions* here is the personal-level grasp of how the world is represented to be by an utterance, so the target is not an over-intellectualized claim about subjects' grasp of the concepts that feature in a semantic theory, like that criticised by Soames, but rather the weaker reading of the claim thus far defended.

Hunter and Fricker are surely right that it is possible for a subject to have a perceptual experience without believing the world to be as it is represented in experience, and that in this sense states of perception need not involve states of belief or knowledge. However,

¹¹ The example is adapted from Hunter 1998: 572

we need to distinguish between a subject's knowledge of *how* the world is represented as being in experience and the knowledge that things *are* as they are represented in experienced: lacking the latter does not entail lacking the former. The subject who is alert to the possibility of machine-induced hallucination knows that the world is represented to him as being such that there is an orange before him even if he cannot, in the particular circumstances in which he finds himself, know that things are this way. Indeed, the very fact that a subject has doubts over his perceptual experience seems to presuppose some grasp of how his experience represents things to be; he must know what the world must be like if his experience is veridical in order to question whether *that* is in fact how the world is.

This point is not made to insist that states of perception must be states of knowledge or belief – presumably there may be people who lack the conceptual resources necessary to form the corresponding belief about their experience. But understanding is by its very nature something that demands conceptual resources of a subject. Once we have this distinction between knowing how the world is represented to be and knowing that it is that way, it is difficult to see how the kind of scenario envisaged by Hunter and Fricker should tell against the idea that understanding involves knowledge of truth-conditions in the relevant sense. To see this, let us distinguish three different – though not mutually exclusive – theses concerning what we might mean in saying that understanding an utterance involves knowing what is expressed by a speaker:

- (i) In understanding an utterance, a subject thereby knows how the world is represented to be by his interlocutor.
- (ii) In understanding an utterance, a subject knows that what he took to have been expressed by his interlocutor is in fact what was expressed by his interlocutor.
- (iii) In understanding an utterance, a subject knows that the world is as it represented to be by his interlocutor.¹²

¹² This list is not intended to be exhaustive of the kinds of knowledge a subject might have in understanding an utterance, but is restricted to the claims that might be relevant to the present discussion. We could add, for example, the claim that in understanding an utterance, a subject knows how his interlocutor *intended* to represent the world to be. This would be relevant in those cases where

(i) is the claim that in understanding an utterance a subject must grasp how things must be in order for the utterance to be true. (ii) is a claim about a subject's knowledge about the accuracy of the auditory experience by virtue of which he understands the speaker's utterance, that how things appear to the subject is the way they actually are. (iii) is a claim about the epistemic potential of understanding, that a subject might know the world to be a certain way on the basis of a speaker's utterance, as utilised in instances of knowledge acquired through testimony.¹³ The claims are not unrelated to one another. Both (i) and (ii) would appear to be conditions of (iii). If a subject is to learn that the world is a certain way, he must know how it is represented to be. He must also know what it was that his interlocutor told him, for were he to have reason to doubt the accuracy of his auditory experience, he would be unable to take advantage of whatever warrant he might otherwise have for the content of the speaker's utterance.

Having distinguished these different claims, it should be clear that the argument presented by Hunter and Fricker fails to tell against (i), the idea that understanding involves knowledge of truth-conditions. When the subject in the aforementioned scenario hears the speaker assert that Cameron is Prime Minister, he has an experience with a particular phonological content, representing the speaker to have uttered some certain words. In the circumstances described, he chooses to remain agnostic about whether the words he took to have been uttered by the speaker were in fact the words uttered by the speaker. This suggests that in the particular case, (ii) does not hold. But there is still some way that the subject takes the speaker to have represented the world to be. The claim of (i) is that if he has understood the speaker's utterance, he must know how the world was represented to be. Were it that the speaker uttered words different to those the subject took to have been said, then the subject would simply have *misunderstood* in the sense of

a speaker intends to convey something more than is literally expressed by his utterance, and so would introduce an interest in the extent to which a subject should be responsive to the pragmatically communicated content of an utterance in order to count as understanding. It is an assumption of the present discussion that we may bracket this issue on the grounds that the semantic features of an utterance suffice to represent the world as being some way and that responsiveness to these is sufficient for linguistic understanding.

¹³ On certain accounts of testimonial knowledge, (iii) may seem implausible; if, for example, one believes that a subject should have evidence in favour of someone's reliability in order to be justified in accepting their testimony, then understanding someone's utterance may not be sufficient to have knowledge that the world is a certain way. However, McDowell (1998e) defends a view of testimony that can be interpreted as sympathetic to (iii). See also Evans 1982: 310-11.

having misheard, in which case the antecedent of the conditional would go unfulfilled.¹⁴ To the extent that he counts as having understood the utterance, having heard the speaker's words correctly, the subject hears the world to have been represented a certain way, viz. that Cameron is Prime Minister. Nothing about the situation as described gives the subject reason to doubt his knowledge of how the world is represented to be by these words. In thus knowing how the speaker's utterance represents the world to be, the subject's situation is one that satisfies, rather than tells against, (i).

States of understanding are not the same as the states of perception that give rise to them. One might know oneself to have heard perfectly the string of sounds emitted by a speaker without understanding what was expressed; I might, for example, be certain that I heard someone utter the words 'Schnee ist weiss' and be able to replicate the sounds accordingly, but if I lack knowledge of German then I may not understand that the speaker said that snow is white. Conversely, one could understand an utterance without being sure that one's auditory experience of hearing a speaker is accurate. But insofar as having an understanding-experience is a matter of being confronted by something contentful, a subject must know what that content is. Even if a subject entertains doubts about the accuracy of the auditory experience, worrying that what he took to be expressed may not in fact be what the speaker expressed, this presupposes a grasp of how the world is represented to be by the words he took the speaker to express. That is, he must know how the world is represented to be in order to question whether *that* is what the speaker said by his utterance. To understand, therefore, is to incur a commitment about the meaning of an utterance: it involves knowing how it represents the world to be.

What, then, should we say about the problem Hunter and Fricker take themselves to be addressing? We noted earlier that Hunter and Fricker are interested in the issue of how a subject's experience of understanding an utterance could ground knowledge of what is expressed by the utterance. We saw in the quotation from Hunter that it was supposed that in order for understanding to play the justificatory role for our knowledge of what is expressed by a speaker, it cannot itself involve knowledge of what is expressed. But

¹⁴ There is, of course, another sense of having misunderstood in which the subject hears the words uttered by the speaker perfectly well but attaches the wrong meaning to them. That is not the sense of misunderstanding that arises in this case.

having distinguished between the different interpretations of what we might mean in speaking about *knowledge of what is expressed*, we might now ask: what is the problem supposed to be? If, as the remarks by Hunter and Fricker suggest, the problem is supposed to be one relating to thesis (ii) – how does a subject’s experience of understanding an utterance ground knowledge that what he took to have been said by a speaker is in fact what was said by a speaker? – then it seems that the presupposition behind the question is itself mistaken. For the issue here is one concerning the accuracy of the subject’s perceptual experience. Yet if states of understanding are independent of states of perception, as described above, why should we suppose that understanding an utterance grounds knowledge that one heard correctly? Cases in which a subject hears an utterance perfectly clearly without understanding it seem to undermine the assumption that understanding is responsible for a subject’s knowledge of what is expressed in this sense. So if the problem is supposed to be one relating to (ii), it is not clear there is any problem in need of address.

If, on the other hand, the problem, relates to (i) – how does understanding ground knowledge of how the world is represented to be by an utterance? – then it is unclear that we have been offered any reason to refrain from the obvious reply: understanding an utterance is sufficient to know how it represents the world to be. To understand an utterance is to have an experience in which one is confronted with a content that represents the world to be a certain way. In understanding it, a subject must know how the world is represented to be.

(ii) Distinguishing possibilities

Presented with the argument of the previous section, someone might ask: what is the significance of knowing how the world is represented to be? The answer, I want to suggest, is that knowing how the world is represented to be by a proposition is to know something of the possibilities that are included or excluded – ruled in or ruled out – by the truth of the proposition. When a speaker asserts a proposition, he distinguishes between different ways the world might be. The picture here is one on which there is an array of possible worlds, which together comprise what Wittgenstein in the *Tractatus* dubbed “logical space” (*TLP* 2.202, 3.4, 3.42) – the total collection of possible ways the

world could turn out, or have turned out, to be. The actual world instantiates one among the many ways the world could have turned out to be. Since we are not omniscient, the collection of things we know about the actual world are compatible with numerous other possible worlds in which these things are also true, but which differ in ways as yet unknown to us. So, one way for us to think about our acts of assertion, in which we distinguish between different ways the world might be, is as the effort to narrow down the set of possible worlds we consider to be 'open' given our knowledge about the world. This conception of our conversational practice is one found in Stalnaker, who writes,

To understand what a speaker is doing when she says how things are, we need to understand how she is distinguishing between different ways that things might be. Here I am echoing Dummett, who says that "to grasp the content of an assertion, one needs to know only what possibilities it rules out, or positively expressed, under what conditions it is correct." (Stalnaker 1999a: 2)¹⁵

If we invoke the analogy of language-use as a game between conversational participants, we might think of its aim as being to reduce the set of possible worlds we and our conversational partners consider to be open.¹⁶ On this view, a conversation takes place in a context that determines the set of possible worlds initially considered to be open, according to the background beliefs that the participants suppose to be mutually salient. This defines the parameters of a playing field within which the moves of the game are made. (Things are not always quite so simple, of course, as anyone familiar with philosophical discussion will know; there are occasions when we discover that an interlocutor holds an assumption about which we may be agnostic or which we perhaps reject outright, forcing the parameters within which the conversation takes place to be redefined.) A speaker's utterance will count as informative if it serves to reduce the set of possible worlds theretofore considered to be open. The aim of the game is therefore to move, i.e. to produce utterances, in such a way as to be informative to one's fellow conversational participants.

¹⁵ The quotation is from Dummett's *The Logical Basis of Metaphysics*.

¹⁶ Stalnaker 1999b: 88

Understanding, on this picture, is achieved by a subject if he has some grasp of which possibilities are included or excluded by the proposition expressed by a speaker. So a subject who understands an utterance in English of the sentence ‘Snow is white’ will have some grasp of the fact that the truth of the proposition expressed includes those possibilities in which snow is white and excludes those in which snow is, say, green or black. Likewise, a subject who understands an utterance in English of the sentence ‘Cambridge is north of London’ will have some grasp of the geographical relation between Cambridge and London, knowing that it rules in those possibilities in which Cambridge is north of London and rules out those that are incompatible with this.

I have said that in understanding an utterance a subject will have *some* grasp of the possibilities that are included or excluded by the truth of the proposition expressed. This is because precisely how much a subject should know of how the world is represented to be in order to understand an utterance is partly a matter of the conversational context. Suppose, for example, that my wife decides to recount to me the details of a wedding, informing me ‘The bride’s dress was ivory’. As it happens, the subtle difference in the shades of traditional-looking wedding dresses is lost on me, and I lack the ability to discriminate between ivory, pearl and white. Presented with a patch of pearl or white, I would be just as likely to judge that the dress in question was *that* colour as when presented with a patch of ivory. Yet in this situation, conversational purposes do not require this level of discrimination; I understand to the extent that I know that the dress in question was some shade of marginally off-white (which I would probably identify as white anyway). So there is a sense in which I know how the world was represented to be, even though I am unable to exclude all the possibilities that are strictly excluded by the proposition expressed.¹⁷ However, greater powers of discrimination would be demanded for someone working in a wedding dress shop being told that a customer wanted an ivory dress. In such a context, knowing how the world is represented to be by the utterance

¹⁷ We might contrast this with a humorous scene in Sofia Coppola’s film *Lost in Translation* in which the character played by Bill Murray is sent a range of samples for a new carpet in his study. His wife, knowing that he likes red, has sent various shades, the differences between which are almost imperceptible, with a note that proclaims her approval of the burgundy. The humour of the scene lies in the fact that relative to the parameters of the utterance, i.e. the range of indistinguishable shades before him, the character has no idea which of these the reference to ‘burgundy’ includes or excludes. Although he understands her words (‘I like the burgundy’), in the specific context of utterance he does not know how this represents the world to be in a way that narrows down any of the possibilities initially considered open. There is, then, a real sense in which he does not understand her utterance.

would seem to require grasping that it rules out those possibilities in which the customer wanted a dress in any of the other various shades of off-white.

If something along the lines of this picture is right – if understanding an utterance involves knowing which possibilities are included or excluded by the truth of a proposition – it affords one possible line of explanation in answer to our question of how understanding might explain our knowledge of analytic truths. An analytic truth, being knowable independently of experience, must be one whose constituent concepts or expressions are therefore sufficient to guarantee its truth. That is to say, an analytic truth has truth-conditions which are guaranteed to be fulfilled. However, two clarifications are required. First, although we might rephrase this point to say that the truth-conditions are necessarily fulfilled, it should not be understood as claiming that all analytic truths are necessary truths, in the metaphysically modal sense of being true in all possible worlds. Consider, for example, Kaplan's point about an utterance of 'I am here now'. It seems that one can know this to be true without any investigation into the world, and it looks like its truth-conditions are guaranteed to be fulfilled. But the fact that the speaker is located at a particular place at a specific time is hardly a matter of metaphysical necessity – there are possible worlds in which they might have been elsewhere at that time, or might even have never existed. Second, the idea that concepts or expressions may guarantee the truth of a proposition does not – in and of itself – conflict with the Quinean sentiment that all propositions depend for their truth on how things are in the world: it is at least possible that the arrangement of some concepts *ensures* that the proposition expressed is true even though it is *made* true by how things are in the world. The example of 'I am here now' is a case in point. We might suppose that the constituent expressions, together with their arrangement, suffice to ensure that on any given occasion of utterance, it expresses a truth; yet that it expresses a truth might be so because of how the world is, namely, that the speaker *is* in the particular place denoted by 'here' at the time denoted by 'now'.

If understanding an utterance involves knowing something of the possibilities that are included or excluded by the truth of the proposition expressed, and analytic truths are those whose truth-conditions are guaranteed to be fulfilled, then someone who understands an analytic truth may recognise that it includes all possibilities we consider

and excludes none. Since grasping that a proposition is true in all possible worlds entails grasping that it is true in the actual world, someone who understands an analytic truth may recognise that it is true.

We can illustrate this rather general account by considering a couple of concrete cases. Take the proposition that no object can be both red and green all over at the same time.¹⁸ Understanding it involves understanding the expression 'red', which we may suppose that a subject does if he knows the role it plays in a speaker's representation of how things are in the world. Presented with a paradigmatic instance of red in clear lighting conditions, a subject should be willing to judge 'That is red'; equally, presented with something that is clearly not an instance of red in clear lighting conditions, a speaker should be willing to judge 'That is not red'. Understanding the proposition also involves understanding the expression 'green', and we may suppose that the general conditions for this hold *mutatis mutandis* as for 'red'. Thus, when a subject reflects on what it would be to be presented with a given shade, he can recognise that an experience on which basis it would be correct to judge that 'That is red' is also one in which it would be correct to judge 'That is not green'. The possibilities in which an object can be described as being red all over exclude all those in which it could be described as being green all over, and so the truth-conditions laid down by the proposition that no object can be both red and green all over at the same time are guaranteed to be fulfilled.

Or consider an utterance of 'I am here now'. The expressions 'I', 'here' and 'now' are indexical expressions whose semantic contribution may vary according to the context of utterance in which they are used. To understand them is to have some grasp of the way their use in a particular context determines a semantic value by which the truth of the utterance is to be evaluated. To understand 'I', for instance, is to know that it denotes whoever happens to be the speaker at that time; to understand 'here' is to know that it denotes a speaker-centric location, that is, a location of contextually determined parameters but which at the very least must include that of the speaker; to understand 'now' is to know that it denotes a time of contextually determined length but which at the

¹⁸ This case is discussed in Peacocke 2000: 266-270 and also Cassam 2007: 196. These have significantly influenced the discussion here.

very least must include that contemporaneous with the occasion of utterance by the speaker. Although the semantic values determined by these rules may express a proposition which is not itself necessary, the way in which the rules determine a semantic value suffice to ensure that an utterance of 'I am here now' excludes any possibility of falsehood. A subject who understands the sentence, aware of the rules by which a meaning is assigned, may recognise that relative to the values by which it is to be evaluated, it excludes all possibilities in which it is false and therefore that its truth-conditions are guaranteed to be fulfilled.

(iii) The return of the metaphysical conception

There is, of course, no requirement or expectation that a subject need conceive of what goes on when he reflects on an analytic truth as *recognising that its truth-conditions are fulfilled*. The phenomenology, we said in comparing accounts of analyticity and rational intuition in the introduction, may simply be that of seeming to 'see' just by virtue of understanding that the proposition in question must be true. Yet because of this, we might suppose that there is some further level of explanation to be had. We might ask the further question: how do we know these truths to be compatible with all possible worlds? By virtue of what do we recognise their truth-conditions to be fulfilled? One could at this point insist that there is no further explanation to be had, that we have identified understanding as the source of our knowledge and explained how it could be responsible for this achievement. But the questions take on significance in view of the dialectic from which we started concerning the origin of the a priori. For if all we are prepared to say is that in understanding an analytic truth one is simply able to recognise that it must be true, then it is unclear that we have come much further than the appeals to a faculty of rational intuition that the invocation of analyticity was supposed to supplant. This is not necessarily an intolerable outcome – to some it might be evidence that there is less difference between the two approaches than might be initially supposed. But such a view would overlook the language-centred explanations of analyticity and the way in which these contrast with the idea favoured by proponents of rational intuition that, in exercising such a faculty, we acquire some sort of direct insight into reality, i.e. a knowledge of truths which do not simply reflect some aspect of our linguistic practice. So,

if it is to be distinguished from an account of this latter kind, it is incumbent upon a satisfactory defence of analyticity that it be able to answer these further questions.

It is instructive to consider the remarks by Wittgenstein which so influenced the likes of Carnap and thus shaped modern invocations of the notion of analyticity. In the *Tractatus* we find the suggestion that those statements true in all possible worlds are not genuine propositions at all. Genuine propositions are those portions within logical space that represent the world to be some particular way; they are, in Coffa's words, "tools for conveying information about how things stand; they tell us something about "the world" in such a way that what they say agrees or fails to agree with the facts." (Coffa 1991: 248) Those statements true in all possible worlds, by contrast, say nothing about the world (*TLP* 5.142) because their role is not to discriminate between the world's being any particular way but rather to define the logical space in terms of which the possibilities of the world are represented (3.42). They must be compatible with all possible worlds precisely because they are what Wittgenstein calls "the scaffolding of the world" (6.124), the structure or framework by virtue of which we assess the various possibilities that might be, or might have been, realised. The thought behind this picture is that in considering the various possible ways the world might be, or have been, we must hold some things as fixed across all possible worlds in order for this way of thinking to have any structure at all. And those statements we hold fixed are those that come out as true in all possible worlds. For example, when we try to envisage the possibility of there being an object that is both red and green all over at the same time, we cannot because we recognise that the properties denoted by 'red' and 'green' exclude one another. But what makes their mutual exclusion relevant is the fact that we treat it as a general constraint on the possibilities we admit that it cannot be that both p and $\text{not-}p$ are instantiated. Thus, $\text{not-}(p \ \& \ \text{not-}p)$ is a part of the framework that structures our thinking about the world and, hence, is amongst those statements we typically class as being true in all possible worlds.

The suggestion that presents itself here is that we can appropriate the framework model utilised by Wittgenstein and Carnap in order to explain our recognition that the truth-conditions of an analytic truth must be fulfilled. We can view the practice of a linguistic

community as being defined by a framework which structures our thinking about the world. This framework comprises the propositions that are held fixed and define the moves that can be made within the practice. Since these propositions have the role of delimiting what we consider to be possible, we hold them as being true in all possible worlds. A subject who is sufficiently integrated within the practice as to understand one of these framework propositions is thereby in a position to recognise the status it has within the practice.

An explanation of our knowledge along these lines blurs the commonly accepted distinction between the so-called metaphysical and epistemological conceptions of analyticity. For we saw in the first chapter that what characterises the metaphysical view, of which Carnap is held up as the foremost proponent, is the distinction between those substantive statements of a system that say something about the system's intended domain of objects and those statements that structure the substantive statements and function instead as rules of the system. Having put that to one side and focused in the subsequent chapters on the epistemological conception, which is altogether less concerned with the division between different kinds of statements and more with what is involved in understanding certain expressions, we now find ourselves back to where we started: the image of the framework with the corresponding distinction between the substantive and constitutive sentences of a language. Even if there is an alternative explanation for how such knowledge is possible that avoids falling back on some feature of the metaphysical picture, the fact that one *can* come full-circle in this way suggests that the alleged distinction between the metaphysical and epistemological conceptions is considerably less clear-cut than its advocates suppose.

The idea that the epistemological conception of analyticity might ultimately rest on broad features of the metaphysical conception is not new. In one of the more famous responses to Boghossian's seminal paper in which the distinction was made explicit, Harman made precisely this point.¹⁹ Unfortunately, Harman's claim is based on Boghossian's own misconception of metaphysical 'truth in virtue of meaning', where this is interpreted to

¹⁹ Cassam also expresses doubt over whether the two can be separated; he does not, however, elaborate on his reservations (Cassam 2000: 53).

mean that an analytic truth both (i) makes a factual claim about the world and (ii) makes it the case that the world is that way. Harman claims that if one has knowledge that p that is to be accounted for purely in terms of one's knowledge of the meaning of a sentence S (where S expresses the proposition p), this can only be because one's knowledge of the meaning of S somehow already includes the knowledge that p .²⁰ This is an awkward principle to reflect upon, but the thought is that if one can know the proposition solely by understanding the sentence, then the knowledge is not based on any knowledge of how the world corresponds. But in this case, one's confidence that p is true must derive from the fact that the meaning of S makes it so. In this way, knowing what S means entails knowing that what it expresses must be true.

Harman argues that the only way this could be possible is if the expressions of a language have meaning only as a result of the conventions a subject has for their use, that is, he says, "by my intentions to use my terms in such a way that S is true." (Harman 1996: 393) Labelling this view "linguistic conventionalism", he outlines it thus:

in intending to use my terms in such a way that S is true, given the way I am using my terms, I know directly that S is true, given the way I am using my terms. Furthermore, my belief that S is true, given the way I am using my terms, is in this case (we need to suppose) constituted by my *using* S as a belief, that is, the belief that p . Given the way I am using my terms, in so using S , what I believe is that p . I have an immediate belief that p , not based on evidence, and in this context such a belief counts as knowledge. (Harman 1996: 393-4, italics in original)

Harman goes on to suggest that, on this picture, the epistemological features of analyticity do not detach from the metaphysical, writing,

In this view, then, my intention to use my terms in a certain way (i) makes S true and so (ii) gives me direct knowledge of the truth of S . Part (i) invokes what Boghossian calls a "metaphysical" notion of analyticity – truth by virtue of meaning. Part (ii) invokes what he calls the "epistemological" notion – knowledge of truth by virtue of knowledge of meaning. In this approach, the epistemological notion is not

²⁰ Harman 1996: 393

independent of the metaphysical notion, as Boghossian says it must be. Indeed, the epistemological explanation depends on the metaphysical explanation. (Harman 1996: 394)

For Harman, the suggestion that we cannot isolate an epistemological conception of analyticity that does not invoke the metaphysical conception is taken as a *reductio* against the tenability of the notion of analyticity, since he regards the metaphysical conception's commitment to 'truth in virtue of meaning' as problematic for reasons broadly similar to those offered against the idea by Quine and Boghossian. However, in claiming that the epistemological conception might involve the metaphysical conception, we have not here appealed to 'truth in virtue of meaning' in the sense in which the phrase is used by Quine or Boghossian, nor have we relied upon Harman's linguistic conventionalism. In fact, the account so far offered is one that eschews the linguistic individualism that underpins the conventionalism that Harman supposes to be necessary to defend analyticity. For we do not, as a general rule, get to dictate the conventions surrounding the use of our words, nor do our intentions manipulate their meanings. We are immersed within a linguistic community, the practice of which is something in which we participate, and it is the wider practice of this community that determines the meanings of our words. Although our intentions may come into play insofar as we may have the intention to use our words in conformity with the wider practice of the community, it does not thereby follow that we imbue the words with their meanings by intending that they should mean what they do. But to have a communal practice of this kind requires that there be some sort of structure, the rules adhered to by those who participate in the practice. The suggestion, then, is that analytic truths are those whose truth-conditions are guaranteed to be fulfilled precisely because they form, or derive from, the framework that structures the practice of the linguistic community which uses the relevant expressions in some particular way. A subject who is sufficiently immersed in a practice, who knows how its expressions are used, is in a position to recognise an analytic truth as one that reflects some rule governing the use of its constituent expressions. It is in *this* way, by viewing the practice of a linguistic community as governed by a framework that affords the division between the constitutive from the substantive claims of the language, that we might regard the metaphysical conception of analyticity as relevant to the epistemological explanation.

Now, one might harbour the worry that the problem with an account of this kind is that it fails to preserve the intuition that, in the case of at least some analytic truths, what we have in knowing them is genuine knowledge of the world because the propositions in question are objectively true. Take, for instance, my belief that for any given proposition p , it cannot be the case that both p and not- p are true. This is something I suppose I know about the world, that it is an objective feature of the world about which I have a true belief; I do not suppose that it is a construct of the linguistic practice in which I find myself participating. Yet the worry is that by explaining our knowledge in terms of the practice to which we conform, our knowledge cannot be seen as the product of tracking how things are in the world. Our practice might have been other than it actually is – governed by a different framework comprising a different set of propositions – even if the world were the same as it actually is. So the fact that we act in conformity with a particular practice is no guarantee that the statements held fixed within that system are objective truths about the world.²¹

The defender of analyticity can respond in one of two ways. The first is to deny that we should give any heed to our intuition that framework propositions are really about the world. We might simply insist that since their role is to structure the substantive sentences of the language by which we are able to make claims about the world, they cannot really make any such claim themselves. Appearances to the contrary can be explained away by the distinction we saw articulated by Coffa (p. 48) between how things appear from inside the system and how it would appear if we were to step back, detaching ourselves to see it from afar. From the inside, the framework propositions might look to be just as much about the world as any of the substantive propositions; but once we appreciate the framework structure that governs our practice, we should be able to let go of this illusion.

A second line of response which accommodates the core of the intuition is to maintain that while our knowledge of an analytic truth is explained by our sensitivity to its status as a framework proposition in our practice, the practice might itself be shaped by how things are in the world. A natural language does not emerge fully formed from a vacuum,

²¹ A version of this concern was put to me by Marcus Giaquinto.

as it were, with a structure arbitrarily designed without any reference to how things are in the world. Language is a tool by which we communicate thoughts about the world and, so we might suppose, has developed and evolved in such a way as to facilitate our communicative enterprise. Thus, just as our practice to some extent shapes our thinking about the world, our experience of the world informs our practice. On this view, the fact that we hold fixed the claim that it cannot be that both p and not- p are true is not an arbitrary dogma but rather reflects something of how we find the world to be. So, although we might explain a subject's knowledge of the claim by appeal to his sensitivity towards our practice (the fact that we hold the claim fixed), the practice itself can be more or less sensitive to how things are in the world. Were we to discover something about the world in which respect it was found to conflict with a proposition held fixed, we might alter the framework of our practice.

This latter response echoes Carnap's defence of analyticity, particularly his response to the objections from Quine. The framework for a system is something developed towards some specific end, and it can be altered and adapted in order to better meet this end should the need arise. A natural language that enables us to say things about the world has such a framework and it too can be altered and adapted in order to better serve this purpose. Our investigations in the world may reveal that certain aspects of the language are less than optimal for our communicative needs, in which case we may adjust our practice accordingly. And, as Quine points out, any statement could be held true or amended so long as we make the appropriate adjustments elsewhere in the system – though some changes will be better or worse suited to assisting in our communicative enterprise. But any such alteration will only ever replace one framework for another with a (perhaps only slightly) different set of propositions that structure the practice. The aim, of course, is to have a framework optimally suited to facilitate the communication of thoughts about the world. The propositions that comprise this framework will be those that accord with how we find the world to be.

(iv) Accommodating disagreement on the framework model

Can the foregoing explanation of our knowledge of analytic truths accommodate the kind of defeasibility that has made the notion of *being in a position to know* so central to the

previous chapters? Is it possible on the account offered here to understand an analytic truth and yet not know it? It may appear that the answer to these questions is negative, that we have succeeded only in closing off the kind of defeasibility we have been trying to make room for through the idea of *being in a position to know*. For if understanding a proposition involves knowing its truth-conditions, so that understanding an analytic truth involves knowing that its truth-conditions must be fulfilled, what room can there be for the possibility that someone might understand it and yet not know it? Or, to put the worry another way, if an analytic truth is one that is in some way involved in structuring the practice of a linguistic community, so that a subject who participates in that community and understands the proposition in question is in a position to recognise its status within the practice, how can anyone fail to assent without thereby impugning their grasp of the practice?

We can begin to allay the concern by considering another question: what are we doing when we disagree about analyticity?²² The framework model offers us an answer to this question. Or rather, it offers us two, since there is more than one way for disagreement to arise on this picture. We have said that the framework is the collection of propositions that structure a practice – the rules of the game, as it were. Yet although in the case of certain formal systems we can perhaps set down the rules that best suit our needs, in the way Carnap envisaged for the foundations of different mathematical systems, this is not so when we find ourselves as participants of a natural language where we are part of a much wider linguistic community. There is no place we are likely to find an authoritative statement of the various rules to which we adhere in conforming to the practice of a linguistic community. Our situation is in this sense more akin to the subject envisaged by Dummett who learns to play chess without having the rules of the game articulated to him, simply receiving correction whenever he makes an illegal move, and in this way coming to understand how the game is played.²³ So, if a natural language can be viewed as having a framework, the propositions by which it is comprised are not so much prescribed as discerned by observing the practice as it is carried out by its participants.

²² I owe this way of thinking about the problem to Mike Martin, who at the outset of this project suggested to me that the entire debate surrounding the notion of analyticity could be read as an attempt to answer the question: what are we doing when we disagree about logic?

²³ Dummett 1993: 96

Of course, we *can* attempt to explicate the rules we follow in our practice. And this is where disagreement of the first kind arises. For we might disagree over whether some particular proposition features in the set of those that structure the practice, that is, whether our practice really is such as to be committed to acceptance of the proposition in question. In a case of this sort, the disagreement concerns what our practice actually is, whether it involves holding as fixed some particular claim; accordingly, the evidence that would be cited for one or another view would appeal to what participants in fact *do*. It is disagreement of this kind that seems to lie behind Vann McGee's claim that the inference rule modus ponens is not truth-preserving. We generally assume that the use of the 'if, then' pairing in English denotes an inference of the form $\langle \text{If } p \text{ then } q, p, \text{ therefore } q \rangle$ that is truth-preserving for all instances of its kind. McGee, however, offers the following as a counterexample to this:

Opinion polls taken just before the 1980 election showed the Republican Ronald Reagan decisively ahead of the Democrat Jimmy Carter, with the other Republican in the race, John Anderson, a distant third. Those apprised of the poll results believed, with good reason:

If a Republican wins the election, then if it's not Reagan who wins it will be Anderson.

A Republican will win the race.

Yet they did not have reason to believe:

If it's not Reagan who wins, it will be Anderson. (McGee 1985: 462)

McGee points out that speakers might accept the conditional statement and affirm its antecedent yet nevertheless refrain from affirming the consequent. Moreover, he argues, whereas there are many familiar cases where upon recognising the entailment of an instance of modus ponens a subject withdraws assent from one of the premises in order to avoid the conclusion, that is not the reaction in this case – a subject who rejects the consequent has no such inclination to give up either the conditional statement or its antecedent.

We need not be concerned here with whether McGee is right that the above example is a counterexample to modus ponens or how we should respond to the specific case. What it

serves to illustrate is the way in which we might disagree over how to characterise the rules that structure our practice. We normally suppose that the rules that guide our use of 'if, then' are such as to make the inference rule $\langle \text{If } p \text{ then } q; p, \text{ therefore } q \rangle$ truth-preserving. McGee's example calls this into question by calling to our attention a way that speakers might regard as appropriate to use 'if, then' on which modus ponens is not truth-preserving. If he were right, we would have to change our characterisation of our practice of using 'if, then'. The onus on proponents of the orthodox characterisation, then, is to show that the data presented by the example can either be accommodated or explained away in such a manner as to preserve the standard interpretation of the practice.²⁴

A second sort of disagreement that can arise on the framework model relates to what was said in the previous section regarding the ability of a particular framework to be more or less sensitive to how we find the world to be, more or less suited to our communicative needs in talking about it. A subject may know what the practice of a community is, how speakers use their expressions and the claims they hold fixed which structure their practice, and yet believe the practice to be ill-equipped to represent certain features of our experience. Rather than disagreeing about what our practice is, we might disagree about what our practice ought to be, with one party arguing that our practice should be revised in order to better accommodate the phenomena in question. Graham Priest, for example, who argues that there can be instances in which both p and not- p are true, is under no illusion that the framework that structures our normal practice does not typically admit of this possibility. He argues, however, that we would be better placed to deal with certain phenomena if we were to revise our practice in such a way that it was no longer part of the framework that it cannot be that both p and not- p are true. Some of the cases cited in motivation for this view are semantic and logical puzzles, such as the liar paradox; others concern our ability to represent certain possibilities that could arise. Illustrating the latter, Priest writes,

Suppose that there is a country which has a constitutional parliamentary system of government. And suppose that its constitution contains the following clauses:

²⁴ For a range of the responses offered in reply to McGee's example, see Williamson 2007: 93, fn. 12,13, 15.

In a parliamentary election:

- (1) no person of the female sex shall have the right to vote;
- (2) all property holders shall have the right to vote.

We may also suppose that it is part of common law that women may not legally possess property. As enlightenment creeps over the country, this part of common law is revised to allow women to hold property. We may suppose that a *de facto* right is eventually recognised as a *de jure* one. Inevitably, sooner or later, a woman, whom we will call 'Jan', turns up at a polling booth for a parliamentary election claiming the right to vote on the ground that she is a property holder. A test case ensues. Patently, the law is inconsistent. Jan, it would seem, both does and does not have right to vote in this election. (Priest 2006: 184-5)

Again, we need not be concerned with whether Priest is right or how we might try to respond to his claim. The point is that we can see his strategy as one of trying to motivate a revision of the propositions we hold fixed with a view to accommodating phenomena of which it is judged that the existent framework is less well-placed to accommodate. If an argument of this sort finds some kind of consensus in its favour, it may result in a change of the framework and, therefore, the practice it structures.²⁵ An opponent may resist such a move by arguing either that accommodation is not necessary, perhaps because there is no genuine phenomenon to be accounted for, or by demonstrating that the existent framework has the resources that enable it to accommodate the phenomena just as well as the proposed revision.²⁶

Crucially, both forms of disagreement that arise on the framework model fall into the category of what in the previous chapter we referred to as "high-level theoretical concerns," which by their nature presuppose that a subject knows what the prevalent orthodoxy is and in what respect their view deviates from it. McGee, for instance, knows what speakers take themselves to be doing when they use expressions 'if, then', just as Priest knows that they tend to believe that it cannot be both that p and not- p are true.

²⁵ This is not to say, of course, that any such change would be immediate or that it would affect all participants of the practice. It is to be expected that many proposed revisions would be of a fairly academic nature, which may eventually take hold and find their way 'downstream', as it were, to a wider group.

²⁶ For an example of just such a reply to Priest, see Smiley (1993).

They each have the belief that the speakers are in error, but they know enough about the practice to be able to make the appropriate adjustments so as to know which thoughts speakers take themselves to be expressing when they use the expressions ‘if’ and ‘not’ and, therefore, to count as understanding the propositions expressed by speakers who use them.

So, the fact that a subject has a high-level theoretical concern about some element of a practice does not impugn his grasp of the practice – he still counts as a participant of the practice who understands what other speakers take themselves to express, even though he has some theoretical reason for withholding assent from certain analytic truths. We may suppose that the subject is wrong, that the problem he perceives to warrant his dissent is not genuine, just as Williamson believes of Peter, Stephen, McGee and Priest. But in having come to believe there to be such a problem, a subject is unable to take advantage of the privileged epistemic position in which he is placed by virtue of his participation in the practice. In this respect their epistemic situations are like that of the subject of a veridical perceptual experience who has reason to worry about the possibility of machine-induced hallucination: they have the means for knowledge at their disposal, yet given other things they believe, it is impossible for them to utilise their entitlements. We said before about the perceptual case that if the subject were to realise that the hallucination machine was broken, he would be able to exploit his position and know that things are as they appear to him. Similarly, although the deviant logicians do not experience the phenomenological ‘pull’ that comes with a perceptual experience, since they have defended their views for long enough that they feel no temptation to believe that all vixens are female foxes,²⁷ they know what the normal practice from which they deviate is; thus, if they were to realise their error, we may presume that they would (all other things equal) revert to those norms and on that basis become able to utilise the entitlement that had been at their disposal all along.

(v) Concluding remarks

Here, then, is how understanding an analytic truth is able to put a subject in a position to know it. Understanding an utterance involves knowing its truth-conditions. In the case of an analytic truth, this enables him to recognise that its truth-conditions must be fulfilled.

²⁷ Williamson 2007: 101

This is made possible by there being certain propositions that constitute the framework which structures the linguistic practice in which the subject participates. These framework propositions are held fixed in evaluating the possibilities that could be realised in the world and as such are counted as true in all possible worlds. A subject who is immersed in the practice and who understands a proposition of this kind is in a position to recognise its status as being such that its truth-conditions are fulfilled.

A subject may have a high-level theoretical concern about some aspect of a practice, where despite understanding what the norms of the practice are, the subject perceives there to be reason to dissent from these norms. Such concerns may arise where either (i) the subject believes that attempts to characterise the framework of the practice fail to represent the actual practice accurately, or (ii) the subject believes that the actual framework is in ill accord with how we find the world to be and is therefore in need of alteration. In cases where a subject holds a false belief of either kind, he may be unable to take advantage of the position in which he is placed by his understanding of the practice, finding himself unwilling to assent to those analytic truths which relate to his theoretical concerns. In this way, we can account for the possibility that someone might understand an analytic truth and not know it while preserving the idea that those who know it do so on the basis of their understanding.

Conclusion

We observed in the introduction that the standard story of analyticity delineates three core issues to which a satisfactory defence should offer a unified response: first, how does understanding explain our knowledge of analytic truths?; second, why might we be inclined to view paradigmatic instances of analyticity as peculiar items to count among our knowledge?; and third, how could it be possible to have cases of disagreement where a subject understands an analytic truth and yet fails to know it? Although the third of these has very much been the focal point for this work, I believe that in addressing it a picture has emerged that enables us to offer a response to the other two.

How, if understanding explains our knowledge of analytic truths, is it possible to understand them and not know them? I have argued that the defender of analyticity can appeal to the epistemic relation of *being in a position to know*, where a subject has the resources required in order to know a proposition already at his disposal but is hindered from exploiting this privileged epistemic position by the presence of some false, undercutting belief that undermines the subject's warrant for the proposition in question. I have argued that this parallels the need for the same sense of *being in a position to know* in describing certain features of our perceptual knowledge. By conceiving of understanding an analytic truth as putting a subject in a position to know it, we preserve the intuitive idea that the subject's understanding is the source of his knowledge, while making room for the possibility that someone might not take advantage of his position and so not know it.

The need to accommodate this possibility constrains the conception of understanding that can underwrite an account of analyticity. It cannot be a precondition of understanding a proposition that a subject be willing to assent to it, since there is always the possibility that a subject may have some reason to withhold assent. We can move to a weaker conception according to which a subject counts as understanding by virtue of his participation in the practice of a linguistic community. For Williamson this has the consequence that there can be no analytic truths on the grounds that one's participation in the linguistic community is consistent with dissent from any given proposition;

therefore, he concludes, there cannot be any truths a subject's acceptance of which is explained by the mere fact that he counts as understanding it. I have argued, however, that the communal nature of our linguistic practice still yields a conception of understanding stronger than Williamson supposes. To understand, on this view, is to know what significance the expressions of the language have for other participants in the practice; so, when a speaker expresses a thought by the use of an indicative sentence, to understand the utterance is to know which thought the speaker expressed.

How, then, does understanding in this sense explain our knowledge of analytic truths? I have argued that understanding an utterance involves knowledge of truth-conditions, which in turn involves having some grasp of the possibilities that are included or excluded by the truth of the proposition expressed. Since an analytic truth is one whose truth-conditions are guaranteed to be fulfilled, understanding an analytic truth enables a subject to recognise that it is true in all possible worlds. This, I have suggested, is made possible by there being certain propositions that have the status of structuring the linguistic practice in which the subject participates. These propositions are held fixed as we evaluate the possible ways that the world could be, or could have been, and so come out as true in all possible worlds. A subject who is sufficiently integrated within the practice and who understands an analytic truth is thereby in a position to recognise that it has this status of being held fixed within the practice.

This explanation provides us with a response to the remaining question regarding the fact that people, particularly when they are initially introduced to the idea of analyticity through paradigmatic instances, often regard it as somewhat strange to count these as items among our knowledge about the world. We find this when met by a response suggesting that they have an air of the obvious, or that they simply reflect the meanings of the constituent expressions. On the account I have offered, these are quite natural responses to have because analytic truths constitute the framework that structures our practice. The model of the framework divides the statements of a language into the substantive claims about the domain of objects about which we wish to speak and the constitutive claims that structure the practice. The analytic truths, being part of the framework, fall into the latter category, as claims that structure our talk about the world

and hence do not count among what we perceive to be the substantive claims about the world. I have suggested that this need not mean that there is no sense in which they are about the world, since our framework may be adapted in accordance with how we find the world to be. But in recognising a proposition to reflect some aspect of the rules by virtue of which we say things about the world, it is natural that it might strike someone as rather trivial and, perhaps, not itself about the world.

Little has been said up to this point to address directly the concurrent debate about what status the truths of philosophy have, although this theme lies just below the surface of the historical discussions about analyticity, since we consider many of these truths to be a priori and thus subject to whatever account we offer by way of explaining the a priori. For many of the logical positivists inspired by Wittgenstein, the ability to label the truths of philosophy as 'analytic' allowed them to separate what they perceived to be fruitless and speculative metaphysical wrangling from the real business of investigating the world through scientific means. And so even amongst those for whom the notion of analyticity was respectable, such as Carnap, we see the idea that analytic truths are not deep truths about the world in the way the defender of rational intuition supposes that we have direct insight into the nature of reality; rather, they are part of a framework altered as we see fit in order to best serve the interests of the scientific enquiries that provide us with the substantive claims about the world.

We see a similar concern with the status of philosophical truths amongst those who have opposed analyticity. For Quine, philosophical claims should be continuous with the claims of the natural sciences; consequently, there are no truths that can be held separate from the empirically respectable, substantive claims about the world, and therefore none which are analytic. In Williamson too we find the suggestion that to label the truths of philosophy 'analytic' would be to ignore the substance or explanatory value we typically suppose them to have – that they would be in some way less substantial and less world-

involving than the truths of other disciplines,¹ whereas Williamson supposes that the methodology of philosophy is broadly similar in many respects to that of certain sciences.²

This is hardly the point at which to begin making any great claims about the philosophy of philosophy. It will suffice to note that on the account presented here, we need not regard analytic truths as being quite so trivial or insubstantial as someone like Williamson supposes. Even if we view philosophy as aiming, at least in part, to draw out or clarify the commitments of our practice and the framework that structures it, we may disagree about what the commitments of our practice are, requiring us to consider different aspects of our practice and our intuitions as speakers in order to reach some sort of conclusion. But we might, as we saw with Graham Priest, also view philosophy as aiming to assess whether our actual practice optimally reflects our experience of the world, advocating adjustments to the framework if there is a mismatch between the two. Again, this introduces scope for considerable disagreement as to whether our framework needs to accommodate some phenomena or whether it already has the resources to do so. So, for the truths of philosophy to be analytic would remove them neither from controversy nor from reference to our experience of the world.

Finally, it should be noted that the arguments presented here do not purport to achieve a complete vindication of analyticity. I started from the intuition that there appear to be truths we know on the basis of our understanding them and subsequently proceeded to show how we might construct an account that preserves this intuition, based around the claim that in doing so we are able to accommodate the possibility that someone could nevertheless understand an analytic truth and still fail to know it. In doing so, I have made certain suppositions about the nature of understanding and our linguistic practice that, while plausible, are certainly not without their points of contention. Were it to be shown that these could not be sustained, it may have consequences for the proposed account; the arguments here must therefore be understood as conditional upon these assumptions. What has been shown is that we can provide an account of analyticity

¹ See, for example, Williamson 2007: 46-8, 52-4

² Williamson 2007 *passim*, esp. 3-7. The irony is that the disciplines to which Williamson points in claiming this are those of mathematics and logic (6, 47), the truths of which the defender of analyticity already supposes to be analytic.

capable of satisfying the desiderata outlined in the introduction, so even if we ultimately have to reconsider the feasibility of an analytic explanation of the a priori, it is not because accounts of analyticity *per se* are unable to accommodate cases of disagreement.

Bibliography

- Alston, W.P. (1986) 'Quine on Meaning', in L.E. Hahn and P.A. Schilpp, eds. *The Philosophy of W. V. Quine* (La Salle, Open Court)
- Awodey, S. (2007) 'Carnap's Quest For Analyticity', in Friedman and Creath, eds.
- Ayer, A.J. (2001) *Language, Truth and Logic* (London, Penguin Books)
- Bach, K. (2005) 'Context Ex Machina', in Z.G. Szabó, ed. *Semantics Versus Pragmatics* (Oxford Clarendon Press)
- Beebe, H. (2001) 'Transfer of Warrant, Begging the Question and Semantic Externalism', *Philosophical Quarterly* 51: 356-374
- Benacerraf, P. (1973) 'Mathematical Truth', *The Journal of Philosophy* 70: 661-679
- Besson, C. (2009) 'Logical Knowledge and Gettier Cases', *Philosophical Quarterly* 59: 1-19
- Boghossian, P. (1997) 'Analyticity', in B. Hale and C. Wright, eds. *A Companion to the Philosophy of Language* (Oxford, Blackwell)
- _____ (2003a) 'Blind Reasoning', *Proceedings of the Aristotelian Society Supplementary Volume* 77: 225-248
- _____ (2003b) 'Epistemic Analyticity: A Defence', *Grazer Philosophische Studien* 66: 15-35
- Boghossian, P. and Peacocke, C., eds. (2000) *New Essays on the A Priori* (Oxford, Clarendon)
- Bohnert, H.G. (1963) 'Carnap's Theory of Definition and Analyticity', in Schilpp, ed.
- BonJour, L. (1998) *In Defence of Pure Reason* (Cambridge, Cambridge University Press)
- _____ (2001) 'Replies', *Philosophy and Phenomenological Research* 63: 673-698
- Borg, E. (2007) 'Minimalism Versus Contextualism in Semantics', in G. Preyer and G. Peter, eds. *Context-Sensitivity and Semantic Minimalism* (Oxford, Oxford University Press)
- Brewer, B. (1999) *Perception and Reason* (Oxford, Clarendon Press)
- _____ (2000) 'Externalism and A Priori Knowledge', in Boghossian and Peacocke, eds.
- Burge, T. (2003) 'Perceptual Entitlement', *Philosophy and Phenomenological Research* 67: 503-548
- Carnap, R. (1956a) *Meaning and Necessity*, second edition (London, Cambridge University Press)

- _____ (1956b) 'Empiricism, Semantics, and Ontology', in Carnap (1956a)
- _____ (1956c) 'Meaning Postulates', in Carnap (1956a)
- _____ (1963a) 'Intellectual Autobiography', in Schilpp, ed.
- _____ (1963b) 'W.V. Quine on Logical Truth', in Schilpp, ed.
- Carus, A.W. (2007) 'Carnap's Intellectual Development', in Friedman and Creath, eds.
- Cassam, Q. (2000) 'Rationalism, Empiricism, and the A Priori', in Boghossian and Peacocke, eds.
- _____ (2007) *The Possibility of Knowledge* (Oxford, Clarendon)
- Chomsky, N. (1986) *Knowledge of Language: Its Nature, Origin, and Use* (London, Praeger)
- Coffa, J.A. (1991) *The Semantic Tradition From Kant to Carnap* (Cambridge, Cambridge University Press)
- Creath, R., ed. (1990) *Dear Carnap, Dear Van: The Quine-Carnap Correspondence and Related Work* (Oxford, University of California Press)
- _____ (2007) 'Quine's Challenge to Carnap', in Friedman and Creath, eds.
- Davidson (2001) 'Radical Interpretation', in his *Inquiries into Truth and Interpretation* (Oxford Clarendon)
- Dummett, M. (1993) 'What Do I Know When I Know a Language?', in his *The Seas of Language* (Oxford, Clarendon)
- Evans, G. (1982) *The Varieties of Reference* (Oxford, Clarendon)
- Fricker, E. (2003) 'Understanding and Knowledge of What is Said', in A. Barber, ed. *Epistemology of Language* (Oxford, Oxford University Press)
- Friedman, M. (2007) 'Carnap's Revolution in Philosophy', in Friedman and Creath, eds.
- Friedman, M. and R. Creath, eds. (2007) *The Cambridge Companion to Carnap* (Cambridge, Cambridge University Press)
- Garcia-Carpintero, M. and Perez Otero, M. (2009) 'The Conventional and the Analytic', *Philosophy and Phenomenological Research* 78: 239-274
- George, A. (2000) 'On Washing the Fur Without Wetting It: Quine, Carnap, and Analyticity', *Mind* 109: 1-24
- Goldman, A.I. (1988) 'Discrimination and Perceptual Knowledge', in J. Dancy, ed. *Perceptual Knowledge* (Oxford, Oxford University Press)
- Harman, G. (1996) 'Analyticity Regained?', *Nous* 30: 392-400

- Heck, R.G. (1995) 'The Sense of Communication', *Mind* 104: 79-106
- Hume, D. (1996) *An Enquiry Concerning Human Understanding* (Oxford, Oxford University Press)
- Hunter, D. (1998) 'Understanding and Belief', *Philosophy and Phenomenological Research* 58: 559-580
- Kant, I. (1998) *Critique of Pure Reason*, trans. P. Guyer and A.W. Wood (Cambridge, Cambridge University Press)
- Kaplan, D. (1989) 'Demonstratives', in J. Almog, J. Perry, and H. Wettstein, eds. *Themes from Kaplan* (Oxford, Oxford University Press)
- Kripke, S.A. (1980) *Naming and Necessity* (Oxford, Blackwell)
- Margolis, E. and Laurence, S. (2001) 'Boghossian on Analyticity', *Analysis* 61: 293-302
- Martin, M.G.F. (2001) 'Epistemic Openness and Perceptual Defeasibility', *Philosophy and Phenomenological Research* 63: 441-448
- McDowell, J. (1998a) *Meaning, Knowledge, and Reality* (London, Harvard University Press)
- _____ (1998b) *Mind, Value, and Reality* (London, Harvard University Press)
- _____ (1998c) 'Knowledge and the Internal', in McDowell (1998a)
- _____ (1998d) 'Criteria, Defeasibility and Knowledge', in McDowell (1998a)
- _____ (1998e) 'Knowledge by Hearsay', in McDowell (1998a)
- _____ (1998f) 'Wittgenstein on Following a Rule', in McDowell (1998b)
- _____ (1998g) 'The Content of Perceptual Experience', in McDowell (1998b)
- _____ (2002) 'Knowledge and the Internal Revisited', *Philosophy and Phenomenological Research* 64: 97-105
- _____ (2008) 'Avoiding the Myth of the Given', in J. Lindgaard, ed., *John McDowell: Experience, Norm, and Nature* (Oxford, Blackwell)
- McGee, V. (1985) 'A Counterexample to Modus Ponens', *Journal of Philosophy* 82: 462-471
- Nagel, E. and J.R. Newman (1958) *Gödel's Proof* (London, Routledge)
- Peacocke, C. (1986) *Thoughts: An Essay on Content* (Oxford, Blackwell)
- _____ (1993) 'How Are A Priori Truths Possible?', *European Journal of Philosophy* 1: 175-199
- _____ (2000) 'Explaining the A Priori', in Boghossian and Peacocke, eds.

- Pollock, J.L. (1987) *Contemporary Theories of Knowledge* (London, Hutchinson)
- Priest, G. (2006) *In Contradiction: A Study of the Transconsistent*, second edition (Oxford, Clarendon)
- Pritchard, D. (2005) *Epistemic Luck* (Oxford, Oxford University Press)
- Pryor, J. (2000) 'The Skeptic and the Dogmatist', *Nous* 34: 517-549
- Putnam, H. (1975) 'The Meaning of 'Meaning'', in his *Mind, Language and Reality: Philosophical Papers vol. 2* (Cambridge, Cambridge University Press)
- Quine, W.V.O. (1953) 'Two Dogmas of Empiricism', in his *From a Logical Point of View* (London, Harvard University Press)
- _____ (1966a) *The Ways of Paradox and Other Essays*, revised edition (London, Harvard University Press)
- _____ (1966b) 'Truth by Convention', in Quine (1966a)
- _____ (1966c) 'Carnap and Logical Truth', in Quine (1966a)
- Recanati, F. (2004) *Literal Meaning* (Cambridge, Cambridge University Press)
- Ricketts, T. (2007) 'Tolerance and Logicism', in Friedman and Creath, eds.
- Russell, G. (2008) *Truth in Virtue of Meaning* (Oxford, Oxford University Press)
- _____ (2010) 'Analyticity in Externalist Languages', in S. Sawyer, ed. *New Waves in Philosophy of Language* (Basingstoke, Palgrave Macmillan)
- Schilpp, P.A., ed. (1963) *The Philosophy of Rudolf Carnap* (London, Cambridge University Press)
- Shapiro, S. (2000) 'The Status of Logic', in Boghossian and Peacocke, eds.
- Smiley, T. (1993) 'Can Contradictions Be True?', *Proceedings of the Aristotelian Society Supplementary Volume 67*: 17-33
- Smith, B.C. (1992) 'Understanding Language', *Proceedings of the Aristotelian Society* 92: 109-141
- _____ (2006) 'What I Know When I Know a Language', in E. Lepore and B.C. Smith, eds. *The Oxford Handbook of Philosophy of Language* (Oxford, Clarendon Press)
- Soames, S. (1989) 'Semantics and Semantic Competence', *Philosophical Perspectives* 3: 575-596
- _____ (1992) 'Truth, Meaning, and Understanding', *Philosophical Studies* 65: 17-35
- Stalnaker, R. (1999a) *Context and Content* (Oxford, Oxford University Press)
- _____ (1999b) 'Assertion', in Stalnaker (1999a)

- Stanley, J. (2005) *Knowledge and Practical Interests* (Oxford, Clarendon)
- Sullivan, A. (2008) 'Truth in Virtue of Meaning', *Canadian Journal of Philosophy* 38: 373-398
- Sullivan, P.M. (2004) "The general propositional form is a variable", *Mind* 133: 43-56
- Wiggins, D. (1997) 'Languages as Social Objects', *Philosophy* 72: 499-524
- Williamson, T. (2000) *Knowledge and Its Limits* (Oxford, Oxford University Press)
- _____ (2003) 'Understanding and Inference', *Proceedings of the Aristotelian Society Supplementary Volume 77*: 249-293
- _____ (2006) 'Conceptual Truth', *Proceedings of the Aristotelian Society Supplementary Volume 80*: 1-41
- _____ (2007) *The Philosophy of Philosophy* (Oxford, Blackwell)
- Wittgenstein, L. (1921/1961) *Tractatus Logico-Philosophicus*, trans. D.F. Pears and B.F. McGuinness (London, Routledge and Kegan Paul)
- _____ (1953/2003) *Philosophical Investigations*, trans. G.E.M. Anscombe (Oxford, Blackwell)
- Wright, C. (2002) '(Anti-)Sceptics Simple and Subtle': G.E. Moore and John McDowell', *Philosophy and Phenomenological Research* 65: 330-348