Rethinking Analyticity:
A defence of philosophy as a conceptual discipline

This thesis is submitted in fulfilment of a PhD in Philosophy at the University

of Kent

Candidate name: James Angove

Word Count: 95,331 (excl. footnotes)

Final date of submission: 06th May, 2018

Acknowledgements

Contents

Introduction

This thesis argues that philosophy is a conceptual discipline. Prima facie, this is an intuitive point. It may even seem obvious that philosophy is not a science, which we might (crudely) think deals with nature more than concepts; or that it is not a form of mathematics, which, say, deals with mathematical relationships between numbers. For philosophy has no interest in testing and prediction,[1] and try as they might, philosophers struggle to eradicate their problems by deductive calculation. For an example case, consider contemporary epistemology's foundational problem. How does anyone know that the world external to themselves really exists—at least, exists outside of their perceptions?

One of the reasons Descartes, in his *Meditations* (1641), made this issue so compelling is that he systematically stripped us of the typical ways of answering empirical and even mathematical questions. That was the point of *methodological* scepticism: were we to justify belief in the outside world by appeal to seeing trees and birds, feeling the wind, and smelling the ocean, Descartes would simply reply, "But can you place absolute faith in these perceptions? Have they not been wrong before?" Of course they have, as when we mistake tress for people in the distance, even though our senses are generally reliable. Consequently, we can find ourselves wondering (at Descartes's urging) whether we can ever be certain that the world outside of perception exists. Given that the route to this predicament was, in a sense, reflective, it should strike us as perfectly natural to surmise that overcoming the problem involves some reflective exercise too—a meditation, perhaps, in Descartes's sense. Thus, we come to one way in which to characterise the philosophical task before us: we must use our reasoning faculties to *solve* the problem, somewhat perhaps as Euclidean geometers or arithmeticians do, although with much less precise tools and under a constant, scrutinising gaze.

Descartes's adopted mode of (apparently) universal doubt—a doubt of everything that is not absolutely certain—is representative of philosophical problematics more generally. To be sure, philosophy does not always subject its readers to that same standard of doubt, but it commonly has that flavour: a feeling that tried and tested empirical methods are of no use to solving distinctively philosophical conundrums. That is expected, for we did not encounter the sceptical problem above as we might discover some curious empirical phenomenon; we did not encounter some object, some living being, or some seeming force and wonder what its nature could be, how it came into existence, how it manages to survive, or what else it affects. And so perhaps

---

[1] The contemporary trend for experimental philosophy, or "X-Phi", aside—see below.

philosophy's characteristic puzzles may not be forthcoming of themselves. Even at their first showing, some effort is required on the part of the philosophical audience to engage with what is supposed to be puzzling. Indeed, the philosophy reader frequently finds herself as though watching a fictional film, where suspending disbelief is partly constitutive of the aesthetic experience. But unlike film, having ruminated on some strange or disconcerting series of thoughts presented to her on philosophy's screen, the reader is easily led down paths that trouble her as a rational creature, for time to come.[2]

Central debates in "metaphilosophy", or the philosophy of philosophy, concern exactly how to characterise all this: what philosophical puzzles really are, what philosophers do to solve them, and whether the route to solution is relevantly similar to the methods of other disciplines, such as (any of) the natural sciences, or mathematics. In this introduction, I will sketch two very different ways of conceiving philosophy's approach (and what it approaches). To survey the field in this respect, we could look at those who work explicitly in some circumscribed sub-discipline of philosophy that philosophises about itself. But although such an area (metaphilosophy) does now have some distinguishable contours, philosophy has always been a self-reflective enterprise.[3] At minimum, throughout many moments of its history, philosophy's practitioners have felt it apt not only to argue about some specified topics but to argue the method used in the process. Appropriately, then, given the sceptical opening, to find out philosophy's conception of itself we shall look at answers to the question of what knowledge is.

A. J. Ayer's classic response to the so-called problem of knowledge (1956)—i.e., how knowledge is possible—begins its opening chapter with a section titled "The method of philosophy" (1956, 7). Ayer asks himself what it is that philosophers are trying to discover when they pose questions such as "What is mind?" or "What is the nature of belief?". In response, he suggests part of the answer: "already knowing the use of certain expressions, [philosophers] are seeking to give an analysis of their meaning" (1956, 8). That is, anyone approaching the question about belief likely has a sufficient grasp of how "to believe" is correctly used in an indefinite number of circumstances. And while this may be prerequisite to answering such a question (1956, 9), a philosopher is responding to an abstract enquiry—What is *belief*? In this exercise, Ayer thinks,

---

[2] I am not being unkind to film: here, the thought-provoking aspect of films is not the analogous element. We suspend disbelief in the *reality* of films, much as we (if receptive) suspend disbelief in the *urgency* of a philosophical problem.

[3] Cf. Williamson (2007, ix): "The philosophy of philosophy is automatically part of philosophy, just as the philosophy of anything else is".

concrete uses do not of themselves suggest an answer.[4] Even so, he argues, philosophers ought to survey the different uses of the words (which signify the topics) they are interested in. So, on Ayer's account, we should investigate whether we have one kind of phenomenon in view when a person says "I believe..." in different ways, and should fairly conclude that we do not. (Compare "I believe in you" with "I believe the Earth is flat".[5]) And then when we pose ourselves questions about belief—say, whether it is possible to believe something false, or something inconceivable, like two and two's making five—we then must begin by appeal to the ordinary concrete uses of relevant terms:

> The argument therefore depends upon certain considerations of language; in the present instance, upon the ways in which we use, or propose to use, the verb "to know". But this does not mean that it is an argument about words, in any trivial sense, or that it is especially tied to the English language. We are concerned with the work that the word "know" does, not with the fact that it is this particular word that does it. (Ayer 1956, 28)

Ayer's example is about knowledge, not belief, but the general point applies: language is to be seen as an entry point, and something of a litmus test, for philosophical enquiry. Speakers may use different languages but share a tendency to use certain words in similar ways, which makes translation possible; in Ayer's view, then, there is something that extends beyond the surface-level use of words: "For our enquiry into the use of words can equally be regarded as an enquiry into the nature of the facts which they describe" (1956, 29).

Now, in the course of this enquiry into the nature of some facts via the use of language, we are bound to come upon stubborn collisions in the way of our reasoning. For example, we grant that one can believe something false, and that two and two's making five is false; that implies we can believe that two and two add up to five. But it seems as though anyone who believes that two and two make five has not understood "two", "plus", and "five". (Further, is understanding a proposition not normally a precondition for believing it?) These are not Ayer's puzzles, but they are typical philosophical puzzles. Ayer's problem is that of how we can ever know, post-Descartes, that something is the case. Starting with the uses of words, Ayer narrows his focus by imagining

---

[4] Cf. Wittgenstein (1953, §§89–90) and Ryle (1962, 454) for some connected thoughts, even though, as Sandis shows (2010), Ayer's approach to concepts is ultimately very different.

[5] As will become clear in this thesis, even here the comparison is superficially simple. "I believe in you" may be words of encouragement; a signal of trust to deliver on promises (as a loyal voter may say, "I believe in Trump!"); or an uncertain mix of both. Similar subtleties are available for "I believe the Earth is flat", on reflection. See my Chapter 3, §2.3 and Chapter 5, §4 for more full explorations of this linguistic nuance.

characters in different scenarios with claims to know that something is the case. We are invited to judge whether or not they really do know—whether or not the subject's use of "to know" is legitimate. The answer he winds the reader towards is that a subject knows that something is the case if, and only if, what she knows is true and she is entitled to be sure that it is (1956, 35).

What has Ayer achieved here? (Historically, not much, since the subsequently termed "tripartite" account of knowledge is infamously open to objection.[6]) In setting out philosophy's methodology, Ayer advocates a prime place for looking at language. But his stated interest is in what the troublesome words and statements *mean*, which he separates in principle from the application of language (1956, 35). Moreover, philosophy's investigation into uses of "knows that" and what we ought to say about them in different scenarios yields, in Ayer's view, an analysis of knowledge itself. Hence why his very first sentence is that philosophy's methods distinguish it from "other arts and sciences", not its subject-matter (1956, 1).

Over time, Ayer's above perspectives have become somewhat dated. In epistemology, the tripartite account of knowledge (frequently referred to as a "justified true belief" definition) is somewhat still in business, though perpetually vulnerable to counterexample, while many, such as Timothy Williamson (2000), regard knowledge as unanalysable, taking it instead as a primitive notion. But it is not only Ayer's epistemology which fell out of fashion: the idea that philosophy should investigate language—whether or not as a way of investigating the world—is itself heavily criticised in certain mainstream corners. Indeed, the philosopher to whom most of my arguments in this thesis are directed, Williamson, recently argued that "few philosophical questions are conceptual questions in any distinctive sense" (2007, 3). Williamson's position, accordingly, directly contradicts Ayer's, since Ayer can say that philosophy has to start with uses of words in a range of circumstances, which provides a distinctive sense in which philosophical questions are conceptual.[7]

Williamson's methodological perspective is well represented among the naturalists of epistemology—for example, take Hilary Kornblith's *Knowledge and its Place in Nature* (2002):

> The idea that philosophy consists in, or, at a minimum, must begin with an understanding and investigation of our concepts is, I believe, both natural and very attractive. It is also, I believe, deeply mistaken. On my view, the subject matter of ethics

---

[6] Perhaps Ayer's account is more nuanced than many typical tripartite accounts, however, for Ayer treats the right to be sure of something as relative to context—i.e., that different contexts provide different standards for being warranted to hold a belief—and this anticipates a major trend of epistemology, so-called contextualism.

[7] Ayer is more direct—and less cautious—about the matter in his youthful work (1936, 44). See my Ch. 1, §1.1 for discussion.

is the right and the good, not our concepts of them. The subject matter of philosophy of mind is the mind itself, not our concept of it. And the subject matter of epistemology is knowledge itself, not our concept of knowledge. In this book, I attempt to explain what knowledge is. (Kornblith 2002, 2)

The evident clash here, between Ayer's kind of perspective and Williamson's and Kornblith's (which do not taken together provide an exhaustive range) over how to understand philosophy itself is the main topic of this thesis. I am arguing that philosophy is a conceptual discipline: that its methods are distinctively conceptual, as Williamson would put it. And even if my view of the matter is, as Kornblith concedes, natural and attractive, it is by no means any longer uniformly accepted (as it largely would have been in the anglophone philosophy departments of Ayer's heyday).

Beyond epistemology, Williamson's vision of philosophy, including his critique of conceptual methods, is, I suggest, relatively influential. Consider Gillian Russell's esteem in a review of *The Philosophy of Philosophy*:

one thing that I ought to say about this book at the outset is that it is written by someone well-informed and well-qualified for the task, someone with an awareness of the history of his subject, for the different options that are already out there, and for recent developments that have made some of those old options seem untenable. This is a book on the philosophy of philosophy written by someone who is good at philosophy. (G. Russell 2010, 40)[8]

However, opposing this, there is a trend of work which argues broadly in favour of understanding philosophy as a conceptual discipline. Notably, Frank Jackson (1998) defends a reductive kind of conceptual analysis (in Ayer's fashion), and Paul Boghossian (1997) defends an analyticity account of conceptual grasp. More recently, others have defended armchair philosophy (Levin 2009; Cappelen 2012) and an ordinary language approach (Sandis 2010; Baz 2012b) from criticisms found in the work of experimental philosophers such as Weinberg et al. (2001) and Knobe and Nichols (2007). More current still is Boghossian's (2011) response to Williamson (2007). Boghossian is responsible, in earlier work, for setting up the narrative of metaphilosophy as a debate over analytic truth (about which more below). It is within this latter dialectic that the present

---

[8] Others, such as Hacker (2009) and A. W. Moore (2009), have been less complimentary about Williamson's awareness of the history of philosophy.

thesis participates, while offering a defence and articulation of a methodology similar in large part to that of Gilbert Ryle and later Ludwig Wittgenstein.[9]

In my view, Ryle and Wittgenstein have been poorly understood over time and are frequently mischaracterised in straw man attacks, as is acknowledged also by, e.g., P. M. S. Hacker (2009) and Constantine Sandis (2010, 2011). The metaphilosophy literature does contain defences of their methodology, or a related one, which sometimes gets called "Ordinary Language Philosophy" or "Oxford Linguistic Philosophy" (sometimes Wittgenstein is felt to operate by different methods to the Oxford philosophers). In this vein are books by Oswald Hanfling (2000) or more recently Avner Baz (2012a), who focuses on Wittgenstein and Austin, and also criticises Williamson, *inter alia*. There is some crossover with my arguments here, but it is questionable to what extent the work of either, which is invaluable and insightful, is taken on board by figures such as Williamson.

One reason, I believe, is the dominance of the following argument, a version of which also features in Williamson (2007, 48–54):

1. To say philosophy is a conceptual discipline means that philosophy is centrally concerned with *conceptual truth*.
2. Conceptual truth just means *analytic truth*.
3. The idea of analytic truth is either incoherent, useless, or somehow defective.
And so,
4. Philosophy is not a conceptual discipline.

There is good evidence to suggest that many contemporary analytic philosophers believe (2).[10] Notably, Boghossian (1997) accepts (2), which is no surprise, since to some extent he initiated this recent trend to use analyticity as an explanation of a priori philosophical investigation. As will become clear, with qualifications concerning what conception of analyticity we have in mind, I too accept a version of (2). Further, in my view, (1) is a roughly correct articulation of the matter. However, I also think that (3) is false; that is, we can explain the idea of analytic truth in a way that is useful and coherent. And when that point is conceded, the argument above is halted in its tracks. Indeed, I hope to show the opposite conclusion; i.e., philosophy *is* a conceptual discipline.

In Chapters 2 and 3, I argue that the only way of explaining what putatively analytic sentences are is to understand them as conceptual truths; specifically, they are norms enjoining the proper application of concepts in a range of circumstances. But clearly, before that argument can be made, we must have some grasp of analyticity, or the analytic-synthetic distinction. That is the point of

---

[9] For a worked-out attempt to group Ryle and Wittgenstein together, in their approach and their arguments, see Tanney (2013). Also see Hacker (1996, 168–72).

[10] See the comments section of Brian Leiter's blog entry, ('Peter Hacker Did Not like Timothy Williamson's Book' 2009). Gillian Russell (2010, 40) also agrees with this point in her review of Williamson.

my opening chapter, which discusses the history of the concept and how it has been variously employed. I will say a brief word about the distinction here first, however, to orient the reader, before overviewing the subsequent chapters.

A textbook statement of the analytic-synthetic distinction, while ultimately misleading, is as follows: *analytic* sentences are those which are true solely by virtue of what their component words mean; *synthetic* sentences are true not only on account of their components' meaning but also on account of facts about how things are. Take the (putatively analytic) sentence, "Adults are grown-ups".[11] Plausibly, it is true because what it is to be an adult is to be a grown-up. And so it would seem strange to say that one ever encounters this claim as a "fact": one does not survey adults to find out whether or not they are also grown-up, and then infer the truth or likelihood of the claim. There is thus a prima facie difference between such a sentence and the (putatively synthetic) "Adults are disillusioned", the truth of which one could not hope to ascertain purely by knowing the meaning of its constituent words (and where a survey, if necessary, might help—or might not, but for a much different reason than in the "grown-ups" case).

Now, given such trivial examples, we really ought to wonder whether the analytic-synthetic distinction, seemingly so pedestrian, could underpin something as grand as the conceptual nature of philosophical enquiry. Chapter 1 shows in fact that the notion of analyticity has been used in attempts to provide just such a foundation, by thinkers as diverse as Leibniz, Kant, Frege, and Carnap. But, following W. V. O. Quine's infamous attacks on analyticity (1951, 1962), in which he characterises the notion as an empty article of faith, many today simply view the analytic-synthetic distinction as untenable (i.e., premise (3) above). Now, it is not news to anyone that Quine had Rudolf Carnap's conception(s) in his sights during his critique. But a key point of Chapter 1 is to show that the different conceptions of analyticity employed throughout history are not all equivalent. Indeed, as I show, philosophers set out varying and overlapping criteria for analyticity, putting the notion to work in sometimes rather different projects. Accordingly, the history helps us realise that there is not "a" conception of analyticity—uniform, ahistorical, and easily understood. This is crucial, for it means that criticism of one conception need not carry over to another; not without some work, at any rate.

In Chapters 2 and 3, I argue that we can and should make sense of the idea that conceptual truths are analytic by considering a workable conception. Specifically, in my view, an analytic sentence is one which expresses a *norm*. Take, e.g., "Adults are grown-ups". In the context of its

---

[11] I use this example instead of the dry, industry-standard "Bachelors are unmarried men", as an homage to Grice and Strawson's defence of analyticity (1956, 150) against Quine (1951), in which they discuss a related example.

normative use, this should be understood as permitting us to re-describe adults as "grown-ups" (and not, say, as "toddlers"). On this understanding of "analytic", analytic sentences are not "made true" by the world (or reality, or "how things are", etc.), because they are *not made true at all*. (Though they are, as I argue in Ch. 3, §3.1, "true" on a qualified understanding of what truth in such cases amounts to.)

Chapter 2's focal point is Boghossian's work on analyticity, in which, effectively, he claims that all sentences are made true by the world, or by the facts. Despite this concession to analyticity's typical critics, Boghossian tries to defend the distinction anew by coining a so-called epistemological conception of analyticity. The epistemological conception holds that knowledge of an analytic sentence's truth requires only knowledge of its meaning, and not knowledge of how things are. His purpose is to show that if some sentences are epistemically analytic in such a fashion, then we can explain some a priori truths, such as those of logic; i.e., we could show that they are knowable purely by virtue of meaning. Against Boghossian, I highlight his fatal presupposition that every declarative sentence is descriptive, whether or not normative. To show why it is fatal, I explore a famed Wittgensteinian example (1953, §50). The upshot of this is that without the presupposition, there is no support for Boghossian's pivotal claim that analytic sentences are true by virtue of describing facts.

In the course of Chapter 2's arguments, a larger purpose emerges. It becomes clear that one reason we may be tempted to think sentences expressing norms are actually descriptive is because in practice normativity is a *diffuse* phenomenon: norms stretch out, their authority fades, and they rule over distinct jurisdictions. Accordingly, we can readily conceive in many cases how what is now normative may soon not be, tempting us to think that any given norm is really a description, since what sentences describe need not be the case. Appreciating the diffuseness and context-sensitivity of norms in turn offers up a sturdy platform to criticise Williamson in the following chapter.

Chapter 3's central argument is motivated by Williamson's agreement with Boghossian that no sentence could be made true by its meaning alone. For, like Boghossian, Williamson also presupposes that sentences which function in practice as rules simultaneously function as descriptions of what is the case. I trace the influence of this false presupposition to problems with Williamson's own criticisms of analyticity, including how he mischaracterises conceptual truth as providing, putatively, a different "sense" of truth. Moreover, I argue that Williamson's attack on the epistemic aspect of analytic sentences—i.e., that in knowing them, we are justified solely by their meaning alone—suffers terribly from the same presupposition. For without the ability to

treat some sentences as uniquely normative, Williamson simply cannot explain how linguistic misunderstanding—*or* understanding—between speakers is possible.

Up to the present point of the thesis, my argument has a dual positive and negative character. Predominantly, Chapters 2 and 3 argue against contemporary positions, by simultaneously invoking and arguing for the normativity of analytic truth. But the following chapter takes a more positive approach still.

In Chapter 4, I present an equally pressing reason to conclude that some sentences express norms and not descriptions: analytic truth, I show, is a precondition for the possibility of synthetic, or empirical, truth. That is, it is impossible to describe the world without a background network of norms enjoining what count as legitimate descriptions. This point is not obvious to some, partly because critiques of the analytic-synthetic distinction often misconstrue syntheticity (empirical description) as much as analyticity. Indeed, I argue, with Donald Davidson and John McDowell, that Quine's way of looking at empirical description, evident in his arguments against analyticity, obscures the possibility for our claims about the world to be justified by it (or, if not our claims, given Quine's holism, then the whole of language or theory). Worse, since Quine's picture of the relationship between language and the world is shared by many others, without a coherent alternative to this picture, *all* truth (not just analytic truth) teeters precariously on the cliff-edge.

To resolve this anxiety, I recite McDowell's (1994) argument that our experience justifies descriptions of the world only because our experience itself is, in an important sense, mediated by concepts. Put another way, McDowell shows us that we must possess concepts in order to experience the world in its conceptually fine-grained nature. And this concept-possession ultimately makes the world accessible to us in experience as the sort of thing we can believe and be justified about. Finally, then, the takeaway message from Chapter 4 is that normative practices, in which we learn analytic truths, are what endow us with this ability to experience and thus describe the world truly or falsely. That is why analytic sentences condition the possibility of synthetic ones.

Chapter 5—the final chapter—presents a prima facie worrying argument, which, if sound, would show the direction of this thesis to be misguided. Let me summarise my line of thought to this point, then, so as to make the threat more vivid. In the foregoing chapters, I have defended the idea that some truths (the analytic ones) are conceptual in nature; they are norms, I have said, grounded in practice. Moreover, these norms that partly constitute meaning, or conceptual content, are often implicit practical principles that language users develop and adhere to over time, which is why I have stressed the supportive role they play. That is, one cites a norm to show that

one is warranted to use a word in a particular way, and a shifting but sturdy backdrop of norms (somewhat like chainmail) makes it possible to describe the world truly at all.

The challenge to this whole dialectic is encapsulated by arguments Mark Wilson makes in his book *Wandering Significance* (2006). After investigating concept-use and conceptual change across a broad range of disciplines, Wilson suggests our grasp of concepts is much worse than we think it is—that in important cases, what our concepts in fact signify strays far from what norms enjoining their correct application will ever tell us. Accordingly, Wilson makes it seem as though the normative backdrop I keep appealing to plays no determinative role for conceptual content; rather, "Nature" forces change and our norms are always a step behind.

In Chapter 5, I respond to Wilson by drawing on my former characterisations of normativity: conceptual norms, I remind again, subtly vary with circumstance, and enjoy a diffuse, changing authority. The problems which Wilson sets out should only worry someone with a simplistic or idealised picture of normativity, according to which norms stubbornly refuse to change. Furthermore, I show that the other drivers of conceptual change on which Wilson focuses, such as the obstacles metallurgists face in applying the predicate "is hard" to a range of materials, not only fail to undermine the constitutive role of norms, they *exemplify* that role. Consequently, the wandering concept-uses which Wilson charts for concepts such as "red", "hard", and "weighs"— which previously seemed to threaten my view—provide a fantastic case in point of how normative conceptual know-how responds to shifting circumstances. Building on Chapter 4, what this demonstrates is that norms—analytic truths—adapt and begin to condition ever new ways of describing the world, in tandem even with empirical investigation.

So much concludes my overview of this thesis, having highlighted the key narrative points in a story about analyticity. But none of it will seem especially relevant or pressing if we do not first have some appreciation of why the analytic-synthetic distinction was coined to begin with, and what work could ever be done with it. Therefore, we now turn to the first chapter.

Chapter One
Conceptions of analyticity

One purpose of this chapter is to show how the notion of analyticity, or the analytic-synthetic distinction, has transmuted throughout history. However, I want to show not just its prevalence throughout the past, but its *significance*. It is no coincidence that philosophers as diverse as Kant, Frege, and Carnap structured their philosophical projects around a distinction between, on the one hand, sentences, propositions, or judgements which are termed "analytic" and those, on the other, which are termed "synthetic". This division of truths is fundamental to each of their philosophical aims. Further, a prima facie similar distinction can be found in Hume's later work, while Leibniz, before him, may well have fathered the idea of analyticity in certain of its most well-known respects.

It will be helpful here to give a broad strokes definition of the distinction to find our footing. Recall the examples from the introduction: a sentence such as "(All) adults are grown-ups" is considered true by virtue of its meaning alone, while "(All) adults are disillusioned" is considered true, if true, by virtue not only of meaning, but also the facts. We must not, however, think this a fair or settled way of drawing the distinction, because over and above the significance of the distinction, I also want to stress its *diversity*. Rarely will any definitions that we consider in this chapter be the same. Frege, for instance, most clearly characterises his distinction via an epistemological criterion. By contrast, the notion which Quine attacks specifies a semantical condition for analytic statements. Again, Kant's distinction is conceptual, or perhaps logical. With respect to the difference between conceptions of analyticity, however, we will also come to see that the different points of emphasis come about precisely because of the work the conceptions are employed to do within specific philosophical projects—however we come to think of those projects decades or centuries on.

I take myself to be offering at first a modest claim in this chapter by arguing that conceptions of analyticity take on particular imprints because of how they are employed. Shortly, we will survey a whole variety of conceptions, most of which are peculiar to particular philosophers, and which exhibit the metaphysical, epistemological, conceptual, and logical characteristics they do because they are embedded in particular philosophical projects. It is a stronger claim, which I also endorse, that it is folly to assess the merits of a given notion of analyticity apart from the system or project in which it figures. Indeed, we should be wary of abstracting an analytic-synthetic distinction from a thinker's project, at least insofar as that leads us to think that there exists a uniform or simple analytic-synthetic distinction across projects. We are liable in that case first to warp how the

distinction played out in other projects, and second to leave ourselves open to criticism when employing it (or criticising it) in our own. For such a distinction does not exist in a vacuum. Kant's notion of the synthetic a priori may well be rejected by most of modern-day analytic philosophy, but one wonders how fair the dismissal is if it fails first to get to grips with his notions of intuition, concept, synthesis, and so forth, which provide it with its special utility. (It may or may not be fair to avoid consideration of the synthetic a priori because of the extra cognitive effort required to understand Kant's critical project, but that is another matter altogether.)

It is a stronger claim still, which I am not endorsing, that the only way to understand the analytic-synthetic distinction is by a near-full awareness of the projects in which it is embedded. That is, one can stress the diversity too much. We should recognise that the different characterisations of analyticity are a matter of emphasis, for there is something intuitive in the distinction to begin with, which normally remains in different articulations. Supposing that English-speakers at random were asked whether they could be sure that adults were grown children, whether they ought to check, I should think they would either suspect a joke was at hand or would get rather frustrated with their questioner. The underlying feeling might be that this question is pointless, that it is fundamentally unlike the question, say, whether adults are disillusioned. I do not think that the technical formulations of different philosophers ever stray that far from something with this much common currency, and so a partial grasp of their formulations is achieved with a partial or minimal grasp of their projects.

In light of the resemblance across the different notions to something possibly already latent within our linguistic practices, a pertinent question arises: do technical formulations of the distinction offer anything over and above the commonplace? Can any important philosophical *work* be done by means of them? Well, what is clear, at least, is that great philosophers have made use of the distinction in order to articulate their vision of philosophy—what it does and what its limits are. Indeed, some critics, such as Quine and Williamson, think that by discarding the distinction we can show that philosophy is not a clearly delimited separate field of enquiry to the natural sciences, but is part and parcel of a broadly empirical system of knowledge (albeit, as a more abstract, theoretical component than, say, biology). In my own way, I agree that philosophy's status as a distinctively conceptual mode of enquiry turns on whether we can make sense of the idea that some propositions are true not because they describe some part of empirical reality but because they reflect something affirmative about our conceptual schemes. What I think is the right gloss on this basic idea, however, is something I offer in the next two chapters.

In sum, then, this chapter serves a number of purposes: it stresses the historical prevalence, significance, and diversity of the analytic-synthetic distinction; it demonstrates that the distinction

morphs to meet the demands of unique philosophical projects; and it offers a glimpse of how philosophy's status, in a broader system of enquiry, is related to whether or not the distinction is intelligible and useful. A further, and final, reason for this historical overview is to anticipate, critically, a contemporary discussion in the next chapter. We will there see that Paul Boghossian (1996/1997) claims to have carved analyticity into distinct general conceptions: metaphysical and epistemological. He wants to be rid of the former, subject as it is (he thinks) to Quinean-style criticisms, and embrace the latter, in order to explain the epistemology of logic. It will become increasingly apparent, however, that his attempt fails at the outset: the analytic-synthetic distinction has been cast in a number of ways, which underline varying aspects—here epistemological, there logical, and so on. But it is highly questionable whether any distinction discussed in this chapter can be framed in the "epistemological only" way to which Boghossian aspires. Indeed, next chapter, we see how this attempt is importantly misguided.

At present, however, we must get back to the roots.

§1 History: Leibniz, Hume, and Kant

We begin the history of the analytic-synthetic distinction with Gottfried Leibniz, but the contrast between the analytic and the synthetic certainly predates him. The ancient Greeks, in the manner of Euclid, employed methods of analysis and synthesis within geometry. Working back from some given, presumed end-point, analysis would proceed to uncover prior principles; synthesis, in reverse, could demonstrate that sought-after conclusion from such first principles. This conception of analysis was thus regressive, contrasting with other conceptions also home to that era—e.g., the decompositional or definitional method of breaking down a complex into simpler constituents; or the transformative method, which turned the analysandum into a new form, such as algebraic expression (see Beaney (2016, §2), for a greater summary of these points). Even in its beginnings, then, there was not an especially settled view as to what analysis and synthesis really amount to, and that ambiguity will persist.

Let us consider an example from the Early Modern era. Marin Mersenne had suggested to Descartes[12] that the argument of his *Meditations* would benefit from a geometrical layout, according to which terms are first defined, then axioms provided, and theorems derived. Descartes, in response, distinguished between two modes of demonstration, analytic and synthetic. Descartes gave the name "synthesis" to Mersenne's suggested presentation, which he tentatively labelled a

---

[12] Mersenne organised the second set of objections from which this point is taken, and is normally taken to be their author; see Descartes (1641, 70), especially the editors' footnote.

posteriori; and he described "analysis" as "the true way by which a thing has been discovered methodically, and, as it were, 'a priori'" (1641, 92). Yet, Beaney has argued (2016, §4), Descartes's descriptions of his philosophical methodology show he was actually committed more to a decompositional-compositional model of analysis and synthesis than the regressive-progressive kind suggested by his appeal to Euclidean geometry.

Either way, Descartes's distinction is perhaps unclear to the modern ear. And that is complicated by the fact that at the turn of the nineteenth century, as we will see in §2, the deductive nature of Descartes's synthetic method will conversely become characteristic of a priori, analytic definition (which is also transformational, for its employment of a logical language). We can already see the fluidity of these distinctions, then, and how they can be dependent for their meaning upon philosophical traditions. Given this, what justifies starting at any given historical point, and not further back; why Leibniz and not, say, Descartes? Indeed, Leibniz did not discuss truth within the framework of an explicit analytic-synthetic distinction, so it is not obvious that we should begin a history of that distinction with his work. Nonetheless, there are reasons to start the ball rolling from Leibniz's vantage point, which I will attend to presently.

First, what Leibniz does say about different kinds of truth is without doubt suggestive of later formulations of the distinction, and much more so than Descartes's. Second, commentators are not shy of attributing some such distinction to him anyway (e.g., see Fitch 2008; Hanfling 1981, 76; Rescher 2011), especially since Leibniz did recommend a process of *analysis* to uncover truth, and not just to demonstrate what has previously been discovered, in Descartes's manner. And third, although perhaps it can fairly be said the distinction as applied to judgements was first Kant's (cf. Anderson 2015, n.8), Leibniz's separate distinction between "truths of reason" and "truths of fact" looks prima facie like an incipient version of it.

In the following subsection, however, I remain ambivalent on attributing to Leibniz a proper analytic-synthetic distinction. For it would be a misleading anachronism to find, in the service of understanding Leibniz's philosophical system, an analytic-synthetic *distinction* in Leibniz; aspects of Leibniz's philosophy proscribe anything like syntheticity or empirical truth as such, notwithstanding his later notion of "truths of fact", which is still in a sense on the "analytic" side. Moreover, as we will see, if we must decide on the matter, the formulation of analyticity most naturally attributable to Leibniz has different consequences when applied to, say, Kant. We will see that Kant's conception of analyticity is certainly inspired by Leibniz's work, but it is in a sense a rejection of the remit of Leibniz's analyticity. If nothing else, then, Leibniz provides a useful contrast case—more useful than Descartes, say—for the conceptions of analyticity covered in the following sections, which certainly do take their cue from Kant.

§1.1 Leibniz

Leibniz constructed an impressive philosophical system, in spite of never expressing it in one major corpus of work. As such, readers are confronted by a vast number of writings, including his essays and letters; this presents some interpretive difficulties in general, as it does for me insofar as my own aim here is to embed his views on truth within a wider project. Nonetheless, such was the systematic nature of Leibniz's thought, it is possible to glean his more-or-less definitive views on a range of topics, and to do so by taking any number of entry points into his system. Plausibly, then, even if one seeks to understand his metaphysics or theology, a good place to start is with Leibniz's view of truth. (See Burnham (2017) for a similar exposition rationale.)

Indeed, in his *Discourse on Metaphysics*, Leibniz characterises the notion of being or substance in terms first of truth (1686, §8). The nature of what Leibniz termed "individual substances"—the basic ontological building blocks of the universe—held an enduring and foundational interest for him. There is some interpretational controversy as to whether for Leibniz only "substantial forms", such as human minds or souls (which he later referred to as "monads" (1714)), comprise these substances, or whether he also admits material, corporeal substances into his ontology (see Woolhouse (1998, 10–16) for discussion). We can think of "substantial form", originally an Aristotelian and a scholastic notion, as the manner or shape of something which "gives" that something part of its essence. For example, a trophy is more than the metal of which it is made— it would not be a trophy without, in addition (and at least), the way in which the material is moulded into a distinct form. With respect to living human beings, Leibniz's thought—at least his early thought; see Woolhouse (1998, 10–11)—is that the mind or soul endows corporeal matter with its distinct form, thereby qualifying humans as individual substances, composed of both matter and form (and, as such, as form-matter composites). However, Leibniz's notion of form, here, extends beyond the way in which the form of the trophy makes it what it is; rather, the form of living things determines their "organisation" and "functioning" (Woolhouse 1998, 13).

Now, according to Leibniz, individual substances (be they immaterial or form-matter composites) are what we express by grammatical subjects—i.e., those things to which we can ascribe properties, but which themselves may not *be* properties: "It is certainly true that when several predicates are attributed to the same subject, and this subject is not attributed to any other, it is called an individual substance" (Leibniz 1686, §8). For example, Bobby is tall, round, and balding; since we can describe him by means of these predicates, and since "Bobby" does not in turn describe some other subject, he is a substance, unlike the properties "tall" or "balding", which are accidental and themselves apply to other subjects besides Bobby.

To get a firmer grasp of individual substance, Leibniz thinks we must become clear about what a true attribution really consists in. His answer articulates a key principle, sometimes referred to as the "predicate-in-subject principle" (Mates 1986, 84), on which other parts of his system rely:

> all true predication has some foundation in the nature of things, and when a proposition is not identical, that is to say when the predicate is not expressly included in the subject, it must be virtually included in it. This is what philosophers call *in-esse*, and they say that the predicate *is in* the subject. So the subject term must always involve that of the predicate, in such a way that anyone who understood the subject notion perfectly would also see that the predicate belongs to it. (Leibniz 1686, §8)

It is worth clarifying the intricacies of these few sentences, not least because the subject-in-predicate principle was foundational to the rationalism which Kant targeted—e.g., Anderson notes that Wolff, Baumgarten, and Meier were all committed to versions of the principle (2015, 10). Let us unpack some of this detail then.

First, the idea of a proposition's being "identical" may seem odd to the modern reader, except insofar as one has in mind trivial identity statements. But it is explicable within the framework of predicate and subject terms being, so to speak, contained in each other. There is some ambiguity in Leibniz's writings as to whether this identity is meant to hold between the objects of predicates and subjects or the conceptual terms themselves; Kant, in fact, argues that this is Leibniz's "fundamental mistake" (Kant 1787, §A270/B326; Beck 1978, n.11). To my mind, this ambiguity is at least consistent with the isomorphic nature of Leibniz's metaphysical system: concepts and objects are mirrored at every step.[13] In spite of Kant's criticism, he did take over the containment metaphor (see my §1.2). At the conceptual level, the idea is that concepts or notions may, so to speak, absorb each other, perhaps in the way we may express meaning-relations today by saying "Well, 'warm-blooded' is *part* of what is meant by 'mammal', since nothing can be a mammal which is not also warm-blooded". Let us take "A mammal is warm-blooded" as our example proposition. In the system of Leibniz, propositions are "identical" when the predicate-term is included within the subject-term; accordingly, this identity holds just as much for our example proposition as for "A mammal is an animal" or "A mammal is a mammal"—all are identical.[14]

---

[13] Henceforth, for simplicity, I will treat Leibniz's subject-in-predicate principle as applying to the conceptual terms and not the objects or substances/properties themselves.

[14] There is a stronger sense in which these propositions are identical for Leibniz, which invokes his so-called universal calculus. In nuce, containment-relations can be defined such that "A contains B" = "A = A+B"—accordingly, there is a literal identity in all affirmative propositions for Leibniz.

Second, and most surprisingly, it is not just concepts such as "mammal" and "warm-blooded" which stand in these containment-relations. It is *any* predicate-term which truly describes a subject-term. That is what Leibniz means when he says that a predicate may be "virtually" included in a subject if not expressly so. According to Leibniz, the analysis of contingent propositions is asymptotic; it must be infinite, approaching ever-closer to identity, and only demonstrable by God. Thus, for Leibniz, "Bobby is balding" is a *near*-identical proposition: because Bobby is actually losing his hair, the notion of the predicate, "balding", is, on an infinite analysis, contained within the notion of the subject, "Bobby". At this point, the isomorphic mirroring between notions and metaphysical substances comes to the fore, since "anyone who understood the subject notion"— i.e., anyone who had a complete grasp of the concept, "Bobby", and so knew all there was to know about him as an individual substance—"would also see that the predicate belongs to it" (1686, §8).

Now, as anticipated, it is clearly not plausible for any mere mortal to know enough about Bobby or any other substance to be able to discern all of his properties; for Leibniz, however, this demonstrates our incomplete grasp of concepts, in contrast to an omniscient God. Thus, God has a complete grasp of all concepts, and *complete* concepts present a fully indexed ledger, as it were, of all included properties (including those in the past and present):[15]

> This being so, we can say that the nature of an individual substance or of a complete being is to have a notion so complete that it is sufficient to include, and to allow the deduction of, all the predicates of the subject to which that notion is attributed. (Leibniz 1686, §8)

Here is where more modern ideas of analyticity may seem applicable. In the work of subsequent thinkers, there is a felt urge to carve off some set of statements, judgements, or propositions which by means of some cognitive act alone one can determine to be true; one can, in a word, "deduce". The proposition "Animals are warm-blooded" is an ideal such candidate. In effect and by contrast, however, Leibniz is arguing that *all* propositions may be deduced in like manner, at least in principle, by God. And that is due to the very nature of individual substances, which is "to have a notion so complete" as to allow this deductive process.

---

[15] Actually, there is some difficulty in working out just where the time-indexes fit in, for Leibniz. It is most intuitive, I think, for Leibnizian complete conceptual content to consist in conjoined and indexed descriptions ("Bobby is balding in 2011", etc.). However, Mates argues this is inconsistent with other key aspects of Leibniz's views on truth. (See Mates 1986, 87–89 for discussion; 91f for a workaround.)

This deductive (or, if it is not tendentious, analytic) parity among propositions is fleshed out further by Leibniz's so-called principle of sufficient reason: the demand that everything be intelligible; that all must have some reason or ground for their being so:[16]

> *sufficient reason*, in virtue of which we hold that no fact could ever be true or existent, no statement correct, unless there were sufficient reason why it was thus and not otherwise—even though those reasons will usually not be knowable by us. (Leibniz 1714, §32)

Leibniz's principle of sufficient reason (articulated here in the later *Monadology*) serves to make networks of true claims together intelligible and rational; it is thus that he speaks in terms of "demonstrating" the truth of propositions, even those we are normally inclined to view as empirical, such as "Bobby is balding" or "the ball is red". For if we are, *per impossibile*, like God, in possession of a complete concept of, say, that ball in particular, we can show how its being red follows from its concept. Leibniz thus sees justification as a matter of deduction, and so to the extent that he is committed to analyticity, he presents a logical, perhaps epistemological criterion, which bears that similarity at least with Frege's, as we come to see in §2.1.

Leibniz does attempt to dissuade readers from concluding that, since a complete concept contains logically explanatory networks, all true predicates are thereby logically necessary. That is to say, Leibniz, at least on the face of it, warns against thinking that the logical structure of concepts implies it is inevitable that agents will do what they do (because of what is true, diachronically, about them), thereby providing a definitive case against free will. His warning mentions a logical criterion that we will find soon in Hume and Kant (§§1.2-1.3): a person who demonstrated the truth of "Caesar crossed the Rubicon" would have shown "that it was rational and therefore definite that this would happen, *but not that it is necessary in itself, or that the contrary implies a contradiction*" (Leibniz 1686, §13—my emphases).[17]

To deal more fully with the accusation of determinism, Leibniz attempts to show how there can be many kinds of truth. This topology presents another area where similar distinctions to the analytic-synthetic are made; indeed, perhaps this is the most plausible part of Leibniz's system in which we can find something close to the idea of syntheticity (cf. Beck 1978, 86–88).

> There are also two kinds of *truth*: those of *reasoning*, and those of *fact*. Truths of reasoning are necessary, and their opposite is impossible; those of fact are contingent, and their

---

[16] This formulation is similar to Spinoza's own before Leibniz, in his exposition of Descartes's *Principles*; see Spinoza (1663, 246, eleventh axiom).

[17] Throughout this thesis, emphasis is only marked when added.

opposite is possible. When a truth is necessary, the reason for it can be found by analysis, by resolving it into simpler ideas and truths until we arrive at the basic ones. (Leibniz 1714, §33)

Truths of reason comprise metaphysical, logical, and "eternal" truths (Mates 1986, 105). The logical truths are basic laws, e.g., the law of (non-)contradiction (Leibniz 1714, §35), and metaphysical truths, such as the truths of arithmetic or geometry; these latter, he maintains, are true in all possible worlds, and reduce to the first kind of essential truth (Mates 1986, 72–73).

The other main class of truths, those of existence or fact, resolve into further kinds (Burnham 2017, §6a.): first, absolutely universal truths, the kind which must be true of this universe given that it is the most perfect universe (as Leibniz elsewhere argues)—even a miracle could not violate such truths, and their contraries again imply contradictions (Leibniz 1686, §§6, 13); second, universal-physical truths, such as the laws of physics, which could principally be violated by miraculous events, since our universe *could* be other than God allows it to be (1686, §7); and, lastly, *individual* metaphysical truths—i.e., truths about Caesar and other individual substances—which are deducible (to God only) from complete concepts of individual substances but are quite contingent to those substances as they are.

Whether or not this division of truths is effective at defending against the threat of determinism in Leibniz's system is not our concern, but we can presently note some of its interesting features. One is that the logical criterion noted above, concerning whether the contrary of a proposition implies a contradiction, is employed by Leibniz to work out whether some proposition is necessary or contingent. For the principle of non-contradiction—"in virtue of which we judge to be false anything that involves contradiction, and as true whatever is opposed or contradictory to what is false" (1714, §31)—is essentially true; indeed, even God is constrained by logical truths, according to Leibniz (1686, §6). Accordingly, although similar non-contradiction criteria subsequently appear in both Hume and Kant, the general principle is not linked specifically to determining whether some proposition or judgement is analytic (indeed, in §1.2.1, I argue that Hume lacks a notion of analyticity, and so really his criterion applies to the a priori).

Another point to note concerns again the class of essential truths: the further division of truths into a mathematical sub-class which reduce to a logical sub-class brings to mind the logicist project as pursued by Frege, Russell, and Carnap. Analyticity was the tool with which Frege sought to dig down to logical foundations, as we will see (§2.1), but it would not be without caveat to say that for Leibniz arithmetic was analytic (likewise for geometry, which was famously synthetic for both Kant and Frege though not for the logical empiricists, given the work of Hilbert). At the least, we can say that the way Leibniz carves up his class of essential truths is suggestive of later attempts;

his system is so unlike the logicist systems to follow, since he seems to lack a notion of genuinely empirical truths (given that God possesses complete concepts of substances), that if it is correct to say that Leibniz is committed to analyticity, then we must say he believes in both the analytic contingent and analytic necessary.[18] Thus, again if we are in this game, Leibniz regards, for God, mathematics wholly as analytic and necessary, and the laws of physics, again for God, as analytic and contingent (i.e., God *could have* designed the world's physics otherwise, though he chose an optimal course).

Taking stock, it remains questionable, to my ear, whether it makes sense to describe Leibniz's discussions of truth as compatible with an analytic-synthetic distinction applied to judgements or truths. It seems he entirely lacks a contrast here, some notion of syntheticity or empirical truth. One could argue, indeed, that Kant's novel conception of the task and structure of metaphysics, which we discuss in §1.3, is founded precisely on the idea that synthetic truths are possible—that metaphysics can reach outside the realm of conceptual truth in which rationalists had thought it wholly consisted.[19] Perhaps at most one might say that Leibniz could conceive another view of how truths are sometimes made, such as is encapsulated by syntheticity, but rules it out within his system. On the other hand, one can see Leibniz definitely is committed to a certain picture of analysis. In this connection, one could argue that Leibniz's project of analysis commits him to a conception of analyticity, as evinced, e.g., in his remarks on the role of analysis in determining truth:

> But in the case of contingent truths, even though the predicate is in the subject, this can never be demonstrated of it, nor can the proposition ever be reduced to an equation or identity. Instead, the analysis proceeds to infinity, God alone seeing—not, indeed, the end of the analysis, since it has no end—but the connexion of terms or the inclusion of the predicate in the subject, for he sees whatever is in the series. (Leibniz 1973, 109)[20]

Thus, even though Leibniz is at pains to emphasise that no human could conduct the analysis of contingent truths, these truths are nevertheless subject to analysis—albeit a different, asymptotic kind of analysis. And if it is a sufficient condition of being analytically true that a proposition be analysable (i.e., reducible to identity, in Leibniz's sense, or asymptotically ever-closer to identity),

---

[18] In his discussion of Leibniz on this topic, and with respect to descriptions of individual substances such as persons, Fitch argues that descriptions may be both contingent and analytic for Leibniz (see Fitch 2008, 38).

[19] See Anderson (2015, §1.1f) for discussion.

[20] This text is edited and published by Parkinson in 1973, though I refer to the author as Leibniz; see bibliography.

then there are no non-analytic truths for Leibniz (cf. Hanfling 1981, 76). Moreover, this process of Leibnizian analysis, as Mates (1986, 122–23) discusses, does proceed by successive substitution of *definiens* for *definiendum* until we resolve a proposition into its identity. This is a conception of analysis that has come to be associated—somewhat wrongly, in my view—with Frege (see my §2.1) and which is prominent in Quine's critique of analyticity as it appears in Carnap (see my §3).

Whichever way we decide on this topic, one thing has become abundantly clear in this survey of Leibniz: the analytic-synthetic distinction, or the notion of analyticity, is far from uniform across philosophical usage. The distinction is intricately connected to the specific projects within which it is employed, and it will scarcely be fruitful to appraise a conception of analyticity when abstracted far from its area of employment.

§1.2 Hume

Along with Leibniz, David Hume is often cited as having prefigured the analytic-synthetic distinction. But perhaps even more so than Leibniz, Hume is thought to have characterised the distinction in ways very similar to contemporary attempts, with respect to its epistemological and modal aspects—respectively, as linked to apriority and necessity. As it happens, I express doubts in §1.2.1 that Hume really is committed to an analytic-synthetic distinction. But nevertheless, without question, Hume occupies an influential spot in its history: contemporary articulations of the analytic-synthetic distinction are bound to bear some or even many of the hallmarks of Hume's thought on his own related distinction. Moreover, Hume is partly whom Kant responds to in his critical philosophy, and it is Kant who formulates the analytic-synthetic distinction in a relatively modern guise. (Rightly or not, in fact, Kant is the first to attribute the analytic-synthetic distinction to Hume.)

§1.2.1 Relations of ideas and matters of fact

The discussion of Hume in this section demonstrates a point which recurs throughout the chapter. Hume's philosophical project—his empiricism—shapes the distinction he makes, in the *Enquiry*, between "relations of ideas" and "matters of fact". This distinction divides two types of "objects of human reason":

> All the objects of human reason or enquiry may naturally be divided into two kinds, to wit, *Relations of Ideas*, and *Matters of Fact*. Of the first kind are the sciences of Geometry, Algebra, and Arithmetic, and in short, every affirmation which is either intuitively or demonstratively certain. . . . Propositions of this kind are discoverable by the mere

operation of thought, without dependence on what is anywhere existent in the universe. (Hume 1777, 20)

Relations of ideas, then, are truths obtainable by reflecting on thinking, and of which mathematical truths of different sorts are paradigm. "Matters of fact", on the other hand, are in effect true claims about the empirical world:

Matters of fact . . . are not ascertained in the same manner; nor is our evidence of their truth, however great, of a like nature with the foregoing. The contrary of every matter of fact is still possible, because it can never imply a contradiction, and is conceived by the mind with the same facility and distinctness, as if ever so conformable to reality. (Hume 1777, 21)

To get to grips with Hume's distinction, it will help to fill in the blanks with some examples. A typical relation of ideas, for Hume, is expressed by a proposition such as "Three times five is equal to half of thirty". How do we recognise it to be a relation of ideas, and not a matter of fact? Well, merely by thinking it through. I may just recall the proposition, but my recollection of it is quite unlike, Hume thinks, my recollection of "The sun rose today". I am not, it seems, just remembering two states of affairs. And my recourse to finding out the truth of each proposition will be different: I will not take it on testimony that three times five is equal to half of thirty; neither will I entertain the thought that three times five could *not* equal half of thirty. Likewise, according to Hume, that the sun rose today is not a matter that can be demonstrated in the sense of it being incontrovertibly *proven*; unlike, as we saw, for Leibniz, where God could actually so demonstrate the proposition, even if we cannot. It is something, instead, that can be estimated to have occurred with a high degree of probability. (Indeed, it is just this degree of probability and its falling short of full-blown demonstration over which Hume agonised.) Further, although "The sun rose today" and "It is not the case that the sun rose today" do contradict each other, there is no contradiction in supposing that either one of these could be true while the other is false.

Let us be more precise by asking how Hume thinks we can tell apart propositions exemplifying relations of ideas on the one hand and matters of fact on the other. At first, Hume gives an epistemological criterion to tell between these different objects of thought, for relations of ideas are held to be "either intuitively or demonstratively certain" (Hume 1777, 20). Dicker explains that for Hume a proposition is intuitively certain when it is self-evident, in the sense that it provides its own evidence: to know its truth we simply have to understand what it says (1998, 36f).[21] A demonstratively certain proposition, by contrast, is one that follows from self-evident propositions

---

[21] Cf. Boghossian's "epistemological conception of analyticity", discussed in my Chapter 2.

(1998, 37)—here we might think of more complicated arithmetical sums which can be demonstrated to follow from those basic ones which are intuitively certain. Thus, relations of ideas are unique for the way that they can be discovered to be true, and the evidence we have for matters of fact must go beyond simply the propositions themselves. (In §1.2.1, however, we will see that an alternative way of characterising intuitive and demonstrative certainty—i.e., as subject to a perceptual model of cognition—radically alters the way we look at Hume's epistemological fork.)

Put another way, as Hume does, the discovery of relations of ideas is not dependent on examining what exists. The ontological claim here is not a criterion for distinguishing relations of ideas from matters of fact, but simply serves to explain why there would be epistemological differences between the two objects of reason. Beyond these aspects, it is of further interest that the criterion Hume offers for how to recognise matters of fact is logically modal (or, as I argue in §1.2.1, actually psychological): the contradictory of a matter of fact is equally as intelligible as its assertion. One can thus tell relations of ideas apart from matters of fact by negating them; if the negation is still intelligible, then one is dealing with a matter of fact: "*That the sun will not rise to-morrow* is no less intelligible a proposition, and implies no more contradiction, than the affirmation, *that it will rise*" (Hume 1777, 21). Note, then, that Hume borrows from Leibniz the law of non-contradiction as a criterion to show what relations of ideas consist in, which Leibniz used as a criterion for determining (analytic) necessary truths.

Accordingly, in contemporary terms, we would describe Hume's distinction between relations of ideas and matters of fact as bearing at least logical, epistemological, and perhaps ontological points of emphasis. Perhaps it is indicative of the importance of the distinction, to Hume, that he articulates it by appeal to all these different aspects. Indeed, the distinction does play a key role for Hume, in the context of his philosophical project, by delimiting his field of enquiry. For, as we shall see, it is matters of fact, and their purported justification, on which Hume centres his attention.

In spite of the fact that Hume cashes out the distinction in these different ways, reading Hume, one gets the impression it is its epistemological character which takes centre stage. But does Hume's epistemological characterisation indicate that he thought relations of ideas were a priori? Juhl and Loomis point out (2009, 3) that this would not seem to be consistent with Hume's thoroughgoing empiricism. A look at Hume's empiricism is warranted, then, to adjudicate the issue.

§1.2.2 The empiricist principle

At the beginning of both the *Treatise* (1740) and the *Enquiry* (1777), Hume introduces another distinction, between "impressions" and "ideas", which allows him to state the guiding empiricist principle for all his philosophical reflection. Impressions (of sensation) are those which "arise in the soul originally, from unknown causes" (1740, §1.1.2).[22] Here Hume has in mind those perceptions (in a very broad sense of the word) which strike upon us, with respect to which we are passive: sensations of heat, cold; pain, pleasure; sight, sound, touch; and so on. Except in a few limiting cases (such as in fever, madness, or dreams; see 1740, §1.1.1), Hume maintains that we can always easily tell apart such impressions from *ideas*, a different kind of perception which recall or echo impressions. For example, a paradigm instance of an idea would be a memory of a sensation such as a pain, some heat, a sight, and so on—"the faint images of these in thinking and reasoning" (1740, §1.1.1). As ideas arise in the mind, reminding one of impressions of sensation, further impressions of reflection emerge (reflective in the sense that they are introspectible). So after, say, a perception of pain, a later idea of that pain may creep upon us, and internally, as it were, we shudder in aversion to it; this last perception—the aversion itself—is an impression of sensation.

With this framework of perceptions in place, Hume is able to state his empiricist principle: "ideas are preceded by other more lively perceptions, from which they are derived, and which they represent" (1740, §1.1.1). Hume regards this principle as entirely general. Thus, contra various rationalists (including Leibniz), for Hume, there are no innate ideas. *Any* idea we can conjure ultimately finds its roots in impressions of sensation—in the soil of our experience in the world.

But to come back to the epistemological worry: Hume's avowed commitment to this empiricist principle should lead us to wonder how he squares it with relations of ideas, given that propositions exemplifying relations of ideas are, in Hume's own words, "discoverable by the mere operation of thought" (1777, 25). That is a classic description of the a priori, after all; and as Reinach (1908, 163–66) persuasively argued, Hume evidently believes in something which shares criteria with Kant's a priori (though see Gotterbarn (1974, 278–79) for an argument that Hume's a priori is psychological). Indeed, in the following sub-section, I argue, *inter alia,* that it does make sense to regard Hume's relations of ideas as a priori, since the propositions instantiating such relations are

---

[22] The *Treatise* comprises three books; references here are first to the book, then part, then section.

held to be intuitively or demonstratively certain.[23] There we will hear, crucially, how for Hume knowledge is delivered via *perception*: we "see", at once, that three times five is equal to half of thirty. And so the Humean a priori takes on a perceptual gloss.

Given the empiricist principle, a larger problem lurks for Hume: the principle itself, that all ideas are derived from impressions, is treated by Hume as though an a priori truth. For example, when, on analysis, difficult ideas like the self or substance are found wanting of an originating impression, Hume declares that minds did not really hold those ideas as such; in that way, metaphysical ideas are debunked. Yet were Hume's principle an empirical generalisation, it would be a more adequate procedure to treat these and other cases as at least contenders for falsifying the principle (see Dicker (1998, 11f) for discussion).

The problem with treating the principle as a priori is most vivid when considering Hume's defence of the principle:

> If it happen, from a defect of the organ, that a man is not susceptible of any species of sensation, we always find, that he is as little susceptible of the correspondent ideas. A blind man can form no notion of colours; a deaf man of sounds. (Hume 1777, 20)

Presumably, then, this support for the empiricist principle is empirical. But that would provide no support for an a priori claim—one "discoverable from the mere operation of thought". Likewise, as Dicker argues (1998, 10f), assuming the principle is a priori, then a thought-experiment concerning (say) a congenitally deaf Beethoven who nevertheless claims to know what sound is, writes down his symphonies, and so on, could not be ruled out as a counterexample by claiming that no such man is known to exist. This is a thorny and interesting issue, but we must leave it at present.

Hume himself turned his attention away from relations of ideas. Instead, he focuses on what he takes to be propositions concerning matters of fact; namely, those about causation, free will, self, and substance. Take "The sun will rise tomorrow". It can be seen not to be a relation of ideas because its contradictory is possible: "The sun will not rise tomorrow", unlike "It is not the case that three times five is equal to half of thirty", is intelligible (1777, 21). This is so because matters of fact state how things are in the world. But how things are in the world is not self-evident; at least, not self-evident in the sense that we cannot work out how things are in the world purely by

---

[23] And as Henry Allison points out, since matters of fact (or the *Treatise* equivalent) are for Hume objects of belief or opinion, "we arrive at the seemingly paradoxical conclusion that for Hume, the radical empiricist, all *knowledge* is a priori!" (2008, 77).

thinking. So, granting that "The sun will rise tomorrow" is classed as a matter of fact, and not a relation of ideas, Hume then examines its justification.

Now, unlike relations of ideas, matters of fact are not demonstrable (note the contrast with Leibniz's truths of fact, which are demonstrable by God). So how can we know them with certainty, if at all? There is an interpretative issue here. Does Hume require for all knowledge that it is certain? Kevin Meeker (2007, 228f; 240f, n. 6) argues that Hume, when being strict, reserves the term "knowledge" for truths knowable with certainty, thus for relations of ideas only. (Allison's point in my above footnote coheres with Meeker's here.) Dicker (1998, 196f, n. 1), however, presents textual evidence to suggest Hume was much more open about this, such that knowledge concerning matters of fact, which involves probabilities and not certainties, was definitely possible. This is not an issue on which we must gain a perspective. All that is required here is to note that Hume at least recognised an epistemological difference between two kinds of proposition, which is either narrow in that both can still be known, albeit differently, or wide, in that only one sort can be known.

I take my above discussion to show that Hume's distinction between relations of ideas and matters of fact helps him to articulate his empiricist project, and so first and foremost this demonstrates the significance of the distinction. Notice how Hume takes as a standard the certainty with which we hold true those propositions concerning relations of ideas. He then measures against this standard our epistemological grasp of propositions concerning matters of fact. Accordingly, in combination with his principle about the derivation of all ideas, Hume then investigates our knowledge of matters of fact concerning causality, the self, and so on.

But notice also how, in turn, the demands of Hume's project characterise his relations of ideas-matters of fact distinction. I take this to show that any proper assessment of Hume's distinction ought to acknowledge the place it has in his project. Indeed, we saw that a profitable line of critique against the distinction involves seeing how it squares up with other aspects of Hume's empiricist system. One might wonder further whether Hume's concern is also semantic. For one way of responding to the abovementioned counterexample to the empiricist principle as a priori is to reconstruct Hume's principle as a test for meaning. This is what Dicker proposes to do (1998, 14–17), taking his cue from Hume's own suggestions in the text. According to the principle thus conceived, if no corresponding impression can be found for an idea (say, "self"), from which it is derived, then that idea is meaningless. I leave it to Hume scholarship to decide whether Hume offered a semantic criterion for relations of ideas-matters of fact—but here we should note that this interpretation is of interest given how the logical empiricists are thought to have extended Hume's thought along these semantic lines, as we see in §2.2.

§1.2.3 Is Hume's a priori analytic?

In the spirit of the historicist theme in this chapter—that no philosophical distinction (such as that between analytic and synthetic truth) can be fully appraised without taking into account its employment within a philosophical project—I must flag at this stage a problem with the reading so far presented: some Hume scholars argue it is mistaken to say that Hume ever *had* an analytic-synthetic distinction (Reinach 1908; Suchting 1966; Gotterbarn 1974; Beck 1978; Owen 1999; Allison 2008). How can this be, for have we not already seen the similarity between that distinction and Hume's own between the propositions concerning relations of ideas and matters of fact? Indeed, prima facie, these two distinctions parallel each other. Yet, as these scholars argue, Hume's fork is most plausibly regarded as only an epistemological distinction, and it is failure to consider Hume's work in the *Treatise* which leads readers to project the analytic-synthetic distinction onto the fork of the *Enquiry* (cf. Beck 1978, 84).

To consider this alternative reading, we must have in view a difference between two models of cognition. On the one hand is a discursive model, a chief proponent of which is Kant. As we will see in greater detail in the next section, Kant argues that there are two fundamental kinds of cognition: intuitions, somewhat like sense-data or impressions, and concepts, which the faculty of understanding applies to those intuitions in order for cognition proper to occur. Thereby, a so-called synthesis between the two kinds of cognition occurs in a judgement, which brings the object of perception under a concept, such as "the ball is red". This exposition lacks detail presently, but it serves here as a useful contrast—the take-home message about Kant's model of cognition is that the objects of knowledge and thought are not bare presences before the eye, or before the mind's eye: they are perceptions as synthesised with concepts in judgements.

By contrast, Hume's model of cognition, much like Locke's before it,[24] appears to be perceptual in nature. That is, for Hume, the fundamental object of thought and knowledge is the item present to one in experience—a sense-datum or impression (and, relatedly, an idea; for the thesis that, for Hume, thinking is as such a kind of feeling, see Allison (2008, 17)). For empiricists, perception is the bedrock of knowledge, and only downstream of that does judgement and inference enter the picture; as such, Hume's model is explicitly not discursive, but perceptual (2008, 6). Moreover, as Allison points out, Hume's two basic items of experience on which judgement builds—impressions and ideas—are conceived in imagistic terms: ideas are but pale copies of the more vivid and lively impressions. Thus, thinking itself is, according to Hume, a stringing together of picture-like items, perhaps later expressed into something discursive. Certainly, then, there seems

---

[24] And Descartes's, see Allison (2008, 6).

to be no room for Hume to accommodate the kind of analytic-synthetic distinction which Kant himself has in mind (and which we explore in the next section) because, by Hume's epistemology, perceptual experience already provides intelligible content (impressions) without the need for any kind of synthesis (see also Allison (2008, 10, 332) on Hume's imagistic interpretation of "concept").

Accordingly, with this contrast between perceptual and discursive models of cognition in view, we can see why there is something fishy about ascribing to Hume a belief in the analytic-synthetic distinction—at the very least in the Kantian sense. Hume's epistemology dictates that all knowledge is delivered in experience, or can be traced, reductively, back to experience. Therefore, even putatively analytic judgements (discovered by the mere operation of thought) could only be known for Hume if they were immediately or derivatively apprehended by perception. Now, at first glance, the notion that one could grasp a relation to hold between some ideas, purely by perceiving those ideas, sounds implausible. Yet if it is true that Hume regards the propositions concerning relations of ideas as a priori graspable, and further that such grasp is to be conceived in imagistic terms, it will follow that the discursive distinction between analytic and synthetic judgements finds no foothold in Hume's own.

Admittedly, it is hard to digest this train of thought without some examples. Let us look, then, at Hume's own from the *Treatise*.[25] Allison argues that Hume continues Locke's tendency to "sensualize" the perceptual model of cognition which they both share with, and inherit from, Descartes (Allison 2008, 7). Locke's discussion of intuitive and demonstrative knowledge strongly suggests a perceptual reading. Take intuitive knowledge—the kind manifested by grasping a judgement that one colour excludes another:

> For in this, the Mind is at no pains of proving or examining, but perceives the Truth, as the Eye doth light, only by being directed toward it. Thus the Mind perceives that *White* is not *Black*, That a *Circle* is not a *Triangle*, That *Three* are more than *Two*, and equal to *One* and *Two*. Such kind of Truths, the Mind perceives at the first sight of the *Ideas* together, by bare *Intuition*, without the intervention of any other *Idea*. (Locke 1689, §4.2.1)[26]

---

[25] Allison (2008, 64) argues persuasively that we have good reason to take the *Treatise* as authoritative on these issues, even though in some respects Hume's treatment of philosophical relations in the *Enquiry* is suggestive of an analytic-synthetic distinction. Moreover, Suchting (1966), Gotterbarn (1974), and Beck (1978) each view the *Treatise* as essential for understanding the fork properly as it occurs in the *Enquiry*.

[26] References to Locke are to first book, then chapter, then section. Citation appears also in Allison (2008, 66).

As can be gleaned from this citation, Locke's picture of intuitive knowledge and truth maintains that we actually *perceive* (in the mind's eye—whatever that quite comes to) such judgements as true; we grasp them at once. From here, it might be thought that Locke must at least appeal to a discursive model to account for demonstrative knowledge—for it seems we must link together intuitively known judgements to demonstrate further ones, thereby initiating a discursive process. Yet this is not so, and Locke's treatment of demonstrative knowledge instructs Hume in the same manner, as we see shortly.

Locke's a priori is defined by certainty, where intuitive knowledge is here the gold standard. In order for demonstrative knowledge to sit on the same level, every rung in a demonstration must itself be intuitively certain—i.e., be perceptible (Allison 2008, 66). For example, once we have seen that "white" is not "black", since the ideas are manifestly different, and seen that "black" is not "yellow" (say), we may "prove" (in Locke's terminology) that "white" is not "yellow". And such a proof or demonstration counts as genuine for Locke only insofar as each judgement may be apprehended fully by intuition, by a perception of the ideas.

Finally, a further aspect of Locke's model which undergirds this intuitive certainty is the nature of apprehension itself: what we perceive distinctly (ideas) are immediately intelligible in their own right:

> In so far as the mind has such ideas, it is aware of their nature or intrinsic content and of their difference from other ideas. For example, to have the idea of white is to know that it is the idea of white and, *a fortiori*, not of blue. (Allison 2008, 67)[27]

Thus, we can glean from Locke's take on colour-exclusion that the phenomenon is first and foremost seen (in some sense of "to see") and thereby grasped; it is not, that is, understood via some discursive act. This immediately makes it suspect that for Locke colour-exclusion (and similar) judgements are "analytic" in any recognisable sense. Indeed, it seems to me that Locke's perceptual model disallows colour-exclusion judgements from falling under any familiar conceptions of analyticity. For example (and to foreground some of the conceptions we discuss in this chapter), these judgements are not held to be "true in virtue of meaning", since the ideas are already intelligible as-is, and thus exclude each other; they are not analytic by being reducible to logical truths via a chain of logical deductions—rather, truistic judgements themselves would seem to be true because they report on a perceived, unanalysable truth at bottom; and, finally, they are

---

[27] See also Locke (1689, §4.3.8). This characterisation of what perceptual experience delivers us is a paradigm of the mythical "Given" we discuss later in Chapter 4; it is fraught with problems, and must be avoided via a Kantian solution.

not normative expressions, since (problematically) the ideas themselves already seem to have normative components built into them (as the idea of white is said to exclude the idea of black).[28]

Now, it is one thing to decide thus that Locke works with no notion of analyticity; however, the similarity of models of cognition between Hume and Locke, ensures that this is true of Hume too. For Hume characterises his own a priori categories, *mutatis mutandis*, in much the same way: the species of philosophical relation (as termed in the *Treatise*) comprising "resemblance", "contrariety", "degree in quality", and "proportion in quantity and number". It is later, in the *Enquiry*, that these relations are effectively rebranded as relations of ideas (see Suchting (1966, 50) and Gotterbarn (1974, 281)).

Take resemblance, for Hume the most general of these relations. Whether two ideas, A and B, resemble each other is something supposedly perceptible by the mind's eye (Allison 2008, 77–78; Hume 1740, §1.3.1). Hume's own example of resembling ideas is again between colours; he notes that ideas of blue and green resemble each other more than those of blue and scarlet. Now, it is true that this example of resemblance imports notions of comparison and degree alien to Locke's own example (as Allison points out), but nonetheless, like Locke, Hume contends that by expressing the resemblance between ideas in a proposition we express a specifically perceptible truth—even though, in fact, many such truths are far from easily perceived. (Think of the resemblance between the same shade of colour, in the first instance isolated in a tube, and in the second, surrounded by patches of other colours.[29]) Moreover, in comparing perceived simple ideas, we are comparing something *ex hypothesi* unanalysable: as Suchting points out, this also disqualifies such a proposition from being analytic, at least insofar as a notion of analyticity relies on an attendant process of analysis (1966, 53).

The best chance of applying some notion of analyticity to Hume's philosophical relations comes when considering the relation of contrariety. Restricted to contrast between that which exists and that which does not, on reflection, instances of this a priori relation hardly seem like the sort we could perceive in reality. For, according to Hume, "existence and non-existence destroy each other, and are perfectly incompatible and contrary" (Hume 1740, §1.3.1); as Allison concedes, this does appear to make a logical point and perhaps anticipates Kant's own criterion for analyticity (2008, 80). Nevertheless, it is because the relation of contrariety does not play well with the perceptual model that it ends up seeming out-of-place in Hume's epistemology, rather than becoming an expansion to it.

---

[28] Following from the preceding footnote, it is this normatively self-standing attribute (or rather, the attempt at it) which qualifies Locke's *ideas*, along with Hume's, as falling prey to the "Myth of the Given".

[29] This example, taken from Wilson (2006, chap.7, §x), will reappear in my Chapter 5.

For the remaining philosophical relations to be discussed are by contrast accounted for by the perceptual model of cognition, even though it may be problematic that Hume conceives them within this model. For example, the degrees of quality between a bright light and a dim one are seemingly discerned via perceptual experience, and thus as expressed in a proposition would more neatly fit into the matters of fact or a posteriori class. Yet, considered as a relation that holds between the ideas themselves—between "bright" and "dim", say—the relation of contrariety is more obviously a priori, even though, again, the way in which we discern the relation is by comparing the perceptual *ideas* and not some essentially discursive concepts.

Finally, consider Hume's relation of proportion in quantity or number. This is the only of Hume's a priori relations by means of which we demonstrate, as opposed to intuit, some knowledge. Accordingly, it concerns arithmetical and geometrical propositions, which are, as with both geometrical and arithmetical, either outright perceived ("[we] might at one view observe a superiority or inferiority betwixt any numbers or figures") or derived, as with arithmetical ones, in an "artificial" (i.e., demonstrative) fashion (Hume 1740, §1.3.1). Such demonstration proceeds via chains of intuitively known (that is, perceived) truths, just as in Locke's case.

Beyond these philosophical relations themselves, it is worth noting that, as Donald Gotterbarn argues, even the law of non-contradiction is based, for Hume, "on our ability to conceive or imagine the relata having or not having various relational properties" (1974, 276). That is, even an aspect of relations of ideas most similar to analyticity—the criterial role played by generating a contradiction—is as much *psycho*logical as it is logical. Thus, the point is not that a denial of a relation of ideas expresses something of the logical form "P is not-P", but rather that the denial would not be imaginable or conceivable in the mind (see also Beck (1978, 83)). Indeed, Gotterbarn points out that Hume's notion of contradiction is itself part of a tradition which spells it out in terms of inconceivability (1974, 277–78).

We can conclude some important things from this brief excursion into Hume's, as well as Locke's, epistemology. In both the *Treatise* and the *Enquiry*, Hume sets aside conceptual room for a class of a priori propositions. In the earlier work, he terms them the philosophical relations among which the relations remain the same so long as the ideas remain the same; in the later work, they are "relations of ideas" as opposed to "matters of fact". Further, it is clear that Hume's model of cognition is by and large perceptual: the rudiments of our thinking and the objects of our knowing are perceived, first and foremost—they are impressions delivered to us in experience or ideas later remembered or imagined, derived from impressions (while the term "idea" covers both these kinds for Locke). Accordingly, Hume treats the propositions which concern relations of ideas in a perceptual, not a discursive, manner. Apriority for Hume as for Locke is a matter of

what can be intuitively or demonstratively known, and these forms of access to the world are themselves imagistically conceived. As such, there is a compelling argument that both analyticity and syntheticity, as applied to Locke and Hume, are misplaced anachronisms.

Taking stock, I have given an overview of a distinction found in Hume that somewhat resembles the analytic-synthetic distinction as later articulated, differently, by various thinkers. And I have given some of the theoretical background necessary for understanding what role that distinction plays in Hume's philosophical project. However, in the course of placing the distinction within the context of Hume's empiricism, serious doubt is cast on the idea that it matches up to any familiar analytic-synthetic distinction. Nonetheless, I hope to have shown at least that the character of Hume's fork is directly affected by the role that it plays within the project.[30] Later, Kant introduces the terminology which has stuck with us to the present day. In fact, Kant claims to find the seeds of the distinction in Hume's fork. Even though I think that is a misattribution, the newly termed analytic-synthetic distinction plays a crucial role in Kant's philosophical project. Indeed, the lines drawn first by the analytic and synthetic, and second by the a priori and the a posteriori, mark out the foundations on which Kant builds his entire conception of philosophy.

§1.3 Kant

> How are *a priori* synthetic judgements possible? That metaphysics has hitherto remained in so vacillating a state of uncertainty and contradiction, is entirely due to the fact that this problem, and perhaps even the distinction between analytic and synthetic judgments, has never previously been considered. (Kant 1787, §B19)

In his *Critique of Pure Reason* (1781/1787), Kant coins the analytic-synthetic distinction as applied to judgements, and proceeds to give it centre stage. As seen in the quote above, Kant signals the importance of the distinction in the Introduction to the *Critique*. Prior to that, in the two prefaces to the work, he explains some important details about his critical project. I cannot possibly hope to give the reader anything like a satisfactory description of Kant's project. But it is fortunate that much of Kant's task is phrased in terms of the analytic-synthetic distinction, so that my limitations here will not be as damaging as they could otherwise be. Accordingly, I will do two things: first, broadly explain the context of Kant's critical project; and second, demonstrate the important role

---

[30] For a rather different interpretation of how Hume treats the relation of ideas/matters of fact distinction, see Meeker (2007, 234-40). Meeker argues that far from being wedded to the distinction, Hume develops a thoroughgoing scepticism with respect to our certainty of relations of ideas, such that he anticipates Quine's revisability claim about analytic truths. Whether or not Meeker is right does not affect my basic claim that the distinction must be understood as part of a wider project.

that the analytic-synthetic distinction plays within this project, and thereby also the character it takes on because of the work it is employed to do.

§1.3.1 Kant's critical project

By the time of the *Critique*, metaphysics was in bad shape, at least according to Kant. Partly due to the rise of Newtonian physics and its battles with a struggling Leibnizian rationalism, and partly due to the anti-metaphysical empiricism of Hume, the intellectual atmosphere did not regard metaphysics with the respect it once commanded. Kant himself compares the differing paths of mathematics, physics, logic, and metaphysics, and paints a sorry picture for his contemporary metaphysicians (1787, B vii-xvi).[31] Indeed, in the second preface to the *Critique*, Kant suggests that metaphysics needs to undergo a change similar to that which Copernicus brought about for physics. Just as Copernicus argued we need to see ourselves—the earth—as what revolves, and the sun as what remains static, so Kant recommends we see not that reason conforms to objects, but that objects conform to reason.

What does it mean to say that objects conform to reason? Roughly, that we perceive and understand what we do, not because, via perceptual or intellectual media, objects project themselves straight into our minds, but because a priori we bring to bear concepts[32] on objects— objects conceived as sensory impressions, that is. This runs counter to the empiricist (and intuitive) thought that the world and its objects are waiting ready to be seen and experienced; that all it takes for perception and knowledge is for us to be present to experience objects. It also runs counter to the rationalist view that the world and its objects may be thought and thereby known by individual subjects. By contrast, according to Kant, *we* act on the objects of experience, by having a priori "intuitions" of the *forms* of perception in general—space and time—which do not arise from perception itself. Further, he argues, we synthesise, bring together, the multi-coloured tapestry of experience into something ordered and tangible—something which can be experienced as unified and objective, and to which we can apply a priori concepts such as causation.

Thus, where Hume had effectively presupposed the order and unity of experience to provide him with the distinct perceptions of (say) a billiard ball's hitting another billiard ball, Kant held

---

[31] And, as is famous, Kant regarded Aristotelian logic as in the main being "a closed and completed body of doctrine", which therefore in comparison with the other "sciences" did not *need* to improve (1787, B viii).

[32] Or, to be more faithful to Kantian terminology, both concepts and intuitions, which belong to the faculties of the understanding and sensibility respectively. Intuitions are understood as unmediated perceptions of objects.

that a priori the subject is already synthesising sensory impressions in order to see a ball *as* a ball, and one as hitting another. Moreover, where Hume looked for the sense impression which gave rise to the idea of one billiard ball's causing another to move after hitting it, and came up short, Kant saw that the concept of causality is applied a priori by the subject to make possible such objective judgements as "Billiard ball A causes billiard ball B to move". That is not to say, of course, that the subject just has a customary feeling which gives rise to his notion of causation—that is Hume's conclusion. Rather, while judgements such as "A causes B to move" are first and foremost subjective, in that they describe *A-[and B-]as-it-seems-to-me*,[33] after the subject applies the concept of causality to perception, such judgements actually attain "objective validity", as Kant puts it in the *Prolegomena* (1783, §18). And so, by the application of this concept, according to Kant, the judgement that A causes B is to hold *really* and not only seemingly, as one would incline to talk about the judgement given Hume's response to the problem.

Accordingly, Kant will go on to argue in the exposition which he terms "transcendental" that experience of objects, of reality, is possible only if certain a priori conditions are met. By means of this exposition, and later what he terms a transcendental deduction, Kant presents a number of compelling reasons to think that reality "in itself"—reality as it is unperceived, or in Kantian terms, "noumena"—is not what we meet with in the world of experience. Rather, *appearances* of objects are what we encounter—in Kantian terminology, "phenomena". Properly speaking, according to Kant, our knowledge of reality or the objects of which it is composed refers only to the phenomenal world. Moreover, he holds, we cannot cognise the noumenal world of things-in-themselves that lays beyond the phenomenal world of experience (1787, §A30/B45), and so the very notion of "object" we are to work with, in examining reason, must be "object for us". That is, an object is only an object within our purview insofar as we have so to speak constructed it—applied an a priori concept to it.

Kant wanted to emphasise, however, that though his critical philosophy embodies a kind of idealism, termed itself transcendental, this is not the sort of Berkelian idealism which denies the existence of matter, in some sense "external" to the subject:

> And we indeed, rightly considering objects of sense as mere appearances, confess thereby that they are based upon a thing in itself, though we know not this thing as it is in itself but only know its appearances, viz., the way in which our senses are affected by this unknown something. The understanding therefore, by assuming appearances,

---

[33] I borrow this "seems-to-me" idiom from Stephan Körner (1955, 53, 64), who uses it to explain how Kant derives his pure concepts of the understanding—the categories—from the (traditional) logical forms of judgement.

grants also the existence of things in themselves, and thus far we may say that the representation of such things as are the basis of appearances, consequently of mere beings of the understanding, is not only admissible but unavoidable. (Kant 1783, §32)

Thus, because the things-in-themselves are necessarily supposed to exist, Kant argues that in no way can he be called an idealist in the Berkelian sense. But it still remains the case that, by Kant's contention, the objects of knowledge are things-for-us: objects which have conformed to our reason, via the a priori application of intuitions and concepts.

This brief exposition wraps up the discussion of Kant's critical project. In the following, I describe the role that Kant's analytic-synthetic distinction plays within this project, and consequently what characteristics that distinction takes on by means of its involvement.

§1.3.2 Kant's analytic-synthetic distinction

Before discussing the analytic-synthetic distinction as it appears in the *Critique*, it should be noted that Kant possibly gives a similar gloss on the distinction, not yet so-named, in his pre-critical period. For instance, in his *Attempt to Introduce the Concept of Negative Magnitudes into Philosophy* (1763), Kant distinguishes between two types of relation exhibited between "grounds" and "consequents": those where the consequent is found, by analysis, to be contained in the ("logical") ground because the consequent is identical with, or with part of, the ground; and those where there is no such identification. The second kind, a relation between "real ground" and consequent, puzzles Kant, and looks to take the form of what he would go on in his critical period to call a synthetic judgement: "that, because something is, something else is [to be]" (Kant 1763, 2:203). Kant's worry is significant, for it suggests that the topic of causality was an impetus to his critical thought, and that an incipient analytic-synthetic distinction provided the form in which he expressed both his pre-critical reflections and his central critical problem. (De Pierris and Friedman (2013, §1) conjecture that Kant likely read a German translation of Hume's *Enquiry* in the 1750s to mid-'60s; if so, potentially there is a tangible line of influence to be drawn between Hume's concern with matters of fact purportedly describing causal relations (and thus his distinction between relations of ideas and matters of fact), through Kant's pre-critical worries, leading eventually to Kant's critical project.)

Now, Kant's analytic-synthetic distinction as it appears in the *Critique* is from one angle very much epistemological in character. (However, cf. Anderson (2015, §1.2.3, 1.3.2) for an argument that the epistemological rendering is not what does the heavy lifting.) For according to Kant, an analytic judgement such as "All bodies are extended" is made with an explicative purpose (1787,

§A17/B11). Such a judgement is classified by Kant in his *Jäsche Logic* (1880, §37) as exhibiting a "non-explicit" identity, which, in contrast to tautologous analytic judgements asserting explicit identities (such as "Man is man"), clarify what was hidden in the subject-concept.[34] For all analytic judgements, then, we may say that one's knowledge is not greatened by the judgement. Those judgements which express non-explicit identities, however, lay out more perspicuously something which one may or may not already cognise (for example, by predicating "extension" of "body"). While analytic judgements, for Kant, are at best explicative, synthetic judgements are by contrast *ampliative* (1787, §A7/B11); they increase what is known about the subject (or in Kant's terms, what is known about the concept of the subject). Accordingly, one way of articulating Kant's analytic-synthetic distinction emphasises these epistemological characteristics.

But this epistemological contrast is also paralleled and preceded by a logical contrast. As is well-known, Kant regarded a judgement as analytic when the concept of its predicate is "contained" (even "covertly") in the concept of its subject. This of course recalls Leibniz, though in practice the containment is (deliberately) less expansive in Kant. As before, one can analyse a concept into its constituent components; if any of these components (or the concept itself) are identical with the predicate in the judgement, then that judgement is analytic. So the concept "body in general", Kant tells us, can be analysed into the constituent concepts, "extension", "impenetrability", "figure", etc. (1787, §B12). If, however, the predicate is not identical with one of these constituents that make up the subject-concept, then the judgement is synthetic. And that deliberately differs from Leibniz, for whom an identity or near-identity always holds of every true affirmative proposition. Sometimes Kant does not speak in terms of containment: "Analytic judgments (affirmative) are therefore those in which the connection of the predicate with the subject is thought through identity" (1787, §A7/B11). However, I take it that as with Leibniz the containment metaphor easily applies: a given subject-concept contains constituent concepts, and upon analysis we reveal those sub-concepts; next, we compare the predicate-concept to this list of sub-concepts, and (in a non-explicit judgement) thereby notice the identity that makes the judgement analytic. This is, then, just another way of talking about containment.

Note that here it is easy to slip from using "analytic" to qualify judgements to talking of analysis *qua* cognitive method. That is, the process by which we discern constituent concepts within more complex concepts, as when we discern "unmarried" and "man" within "bachelor", is itself easily cast in terms of analysis. It looks to be a decompositional kind of analysis, which Kant shares with

---

[34] There is perhaps some textual inconsistency here. Burge points out that Kant, in an unfinished manuscript, classifies identities as non-analytic, being themselves "the limits of analysis"; Burge argues that are not synthetic for Kant either (2003a, 388).

Leibniz (given his predicate-in-subject principle). In that respect, as a response to Leibniz and subsequent rationalists, Kant's recognition of a new class of judgement—synthetic, let alone synthetic a priori—reciprocally makes a substantive claim about cognitive process. Namely, that not all truth is arrived at by ascertaining latent containment relations between concepts, but by actively *synthesising* a connection between concepts.

One way that Kant portrays the method of analysis, in elaborating on the containment metaphor, is by use of a logical principle. Some think this in turn provides a more precise criterion.[35] For some judgement, negate it, and after analysis look for a contradiction:

> All bodies are extended.
> It is not the case that all bodies are extended.

Is the second judgment a contradiction? Well, the analysis of "body" yields what is "thought in the concept", namely "figure", "extension", "impenetrability", and so on (1787, §B12). Accordingly, the second judgement is on analysis a contradiction in terms. A proposition is analytic for Kant, therefore, if and only if its denial is a contradiction in terms. Indeed, Kant affirms the principle of contradiction as "being the universal and completely sufficient *principle of all analytic knowledge*" (1787, §A151/B191), though he earlier hints at the principle's importance when introducing the analytic-synthetic distinction. Here he mentions also that operating with the principle will show an analytic judgement to be *necessary*. Synthetic judgements, by contrast, need not be necessary—though as we shall see, Kant maintains that a class of synthetic judgements are. Note again the contrast with Leibniz, who used the law of non-contradiction as a criterion of a proposition's necessity, but not its analyticity. And Hume appears to invoke the same law, yet he gives it a psychological gloss, such that one can determine a proposition to concern relations of ideas if one cannot imagine or conceive its contrary.

Summing up, in an analytic judgement (say, a non-explicit one), as with Leibniz, the predicate is revealed by analysis to be identical to one of the constituents of the subject-concept; by contrast, in a synthetic judgement, such as "All bodies are heavy" (1787, §B11), one synthesises the predicate "heavy" with the subject-concept "body" (and this conception of synthesis is foreign to Leibniz— fatally, in Kant's view). In analysis, one is only drawing connections between concepts which already hold in a relation of containment, but for synthesis, one must go beyond the concepts.

---

[35] Körner argues that the following gloss on Kant's criterion of analyticity is sufficient to dismiss worries that the containment metaphor itself is vague (1955, 22). Moritz Schlick's (1925, 74–75) clear discussion likewise shows how the principle and containment metaphor relate, along with the Vienna Circle's new concern with convention and definition. See James Van Cleve (1999, 18ff) for discussion of the common criticisms of the containment metaphor.

Kant speaks thus of a requirement or condition on the basis of which the understanding synthesises predicate and subject (1787, §A8f); in the next subsection we see that it is experience which plays that synthesising role in the case of empirical (i.e., a posteriori) synthetic judgements. If most of our judgements are synthetic, perhaps we could frame it such: a judgement in general requires some basis on which it is made, and analytic judgements provide a limiting case, since the understanding relies on nothing other than itself—than logic—to make these explicative moves. That is, in an analytic judgement, analysis simply breaks the subject-concept up "into those constituent concepts that have all along been thought in it" (1787, §B11). Indeed, Kant contends (1787, §B12) that grounding an analytic judgement on experience would simply be absurd (Timothy Williamson, the focal point of my Chapter 3, will very much disagree).

§1.3.3 The synthetic a priori

Given that, on Kant's view, one does not ground analytic judgements on experience, he classifies them as a priori. (And as mentioned above, the test of the principle of contradiction shows analytic judgements to be necessary—for Kant, a sufficient condition of apriority.) Yet synthetic judgements, he argues, are not so simple. The intuitive thought, I take it, is that since analytic judgements contrast with synthetic judgements so far in their logical and especially epistemological respects, synthetic judgements should simply be classified as a posteriori. However, famously Kant proposes another class of judgement: the synthetic a priori. According to how Kant has introduced these terms individually, then, a synthetic a priori judgement is a judgement which is ampliative— which can expand our knowledge—but which we do not make having relied on experience. Kant's notion of synthesis is relevant here again: synthesis, he contends, is made possible through some "X" (1787, §A8). In the case of synthetic a posteriori judgements, that "X" is easy enough to see: experience (of an object). It is because I experience, say, bodies to be heavy that I can connect the predicate-concept to the subject-concept in the synthesised empirical judgement, "All bodies are heavy". But it is not immediately clear just what "X" can allow the synthesis of an a priori judgement, which by stipulation is not analytic.

Take the synthetic judgement, "Every event has a cause". This is not analytic, for Kant, because the notion of a cause is not, he thinks, contained within the notion of an event or happening. Kant contends the judgement is a priori, since it is strictly (i.e., fully) universal; it admits of no exception. As Kant earlier explains, inductive judgements arrived at through a posteriori investigation are only lent an "assumed and comparative universality" (1787, §B4). Further, as he also contends, judgements which are a priori are necessary too (necessity and strict universality are, for Kant, criteria for apriority). Accordingly, as the understanding sees that all events are *necessarily* caused, it

must regard "Every event has a cause" as an a priori judgement. Thus, we come to a class of judgments which Kant claims neither Hume nor Leibniz before him could make out: that of the synthetic a priori.

To recall the quote opening this section, Kant thinks answering the question how synthetic a priori judgements are possible is absolutely vital for metaphysics. And as I discussed shortly thereafter, by Kant's reckoning, Hume has no category of a priori synthetic judgements—"Every event has a cause" is not a relation between ideas, since one can conceive its contradictory, and therefore it is not a priori. Being a posteriori, we must look to experience as the ground of its knowledge. Yet, as Hume found out, there is no impression of necessary connection which gives rise to our idea of it. Accordingly, as Hume sees it, "Every event has a cause" has no worldly justification. This sceptical consequence is of course damning to the enterprise of metaphysics, and its status as a science. One could insist, dogmatically, that causes just do exist in nature, but that is no good for metaphysics either. Thus, Kant saw it as crucial for the status of metaphysics as a serious subject that it answer properly how such synthetic judgements were knowable a priori, since, as he took Hume to show, they cannot be known empirically.

Kant's answer to the question how "Every event has a cause" and other synthetic a priori judgements are possible has been somewhat anticipated in my discussion of his critical project. In a nutshell, subjects actively apply categories (pure[36] concepts of the understanding, as well as pure intuitions) to experience, and one such category is causality. That application of the category, as the missing "X", turns a subjective judgement about perception (say, of events constantly conjoined) into judgements regarded as "objectively valid". The category of causality thereby conditions causal synthetic a posteriori judgements. Accordingly, in the synthetic a priori judgement "Every event has a cause", the notions of event and cause are synthesised in the understanding, neither on the basis of experience, which yields synthetic a posteriori judgements, nor on the basis of logic or concepts, which yields analytic (a priori) judgements.

It is time to review the above discussion. Clearly, for Kant, the analytic-synthetic distinction is foundational. If I am right in §§1.1-1.2, neither Hume nor Leibniz has a comparable distinction, in spite of superficial similarities. Perhaps this is unsurprising, since Kant sees himself as breaking with the Leibnizian tradition in allowing there to be synthetic truths at all, and the Humean tradition by allowing synthetic judgements to be classified as knowable a priori. And like Hume, but to a greater extent, Kant's distinction forms the bedrock of subsequent investigation for him:

---

[36] Pure in the sense that the concepts are completely independent of experience; cf. Kant (1787, A xii).

reason must answer, on pain of losing metaphysics for good to dogmatism or scepticism, what makes possible synthetic a priori judgements.

The epistemological character of the distinction is important, since Kant wants to find out how metaphysics is possible, how we can attain *knowledge* of reality in advance of experience. It would not be apt simply to say, however, that Kant has an epistemic conception of analyticity. For one thing, there are convincing reasons (which we cannot explore) to regard the explicative or ampliative function of analytic or synthetic judgements as *following from* their conceptually different natures—see Anderson (2015, §1.2.3). For another, the demands of his project make clear that analysis has its role, and that this is a logical exercise. Indeed, Kant regards the implications of his own transcendental deduction of the categories as analytic propositions (1787, §B138; cf. Körner 1955, 68–69). But given that Kant distinguishes explicit identities (tautologies) from non-explicit ones, we have no reason to think analyticities on his account are trivialities either. What we do see is that given logical analysis has its place in the project, the analytic-synthetic distinction also has a logical emphasis, in terms of the principle of contradiction, or the containment metaphor, along with the epistemological one.

§2 Recent history: Frege and logical empiricism

§2.1 Frege

Frege's definition of analyticity[37] is best understood in the context of his logicism. The aim of Frege's logicist project, which begins with his *Begriffsschrift* (1879), is to show that arithmetic (though not geometry) is reducible to logic; i.e., to show that arithmetical truths can be derived deductively from logical truths. This project thus has an explicitly epistemological aim, in demonstrating arithmetic to be grounded or secured in logic: grasp of an arithmetical equation is explainable by grasp of logic. This aim is not modest, for arithmetical truths are not prima facie logical (cf. Kanterian 2012, 20f). For instance, take an equation such as "7+5=12"—an attempt to articulate its justification might begin by citing that it is a priori, and logical truths are likely thought to be a priori (though, as we shall see shortly, not all for Frege, since logical laws do not require justification). But how much further can such an articulation go? It seems that we almost have to stop short—seven and five *just do* make twelve. It is not clear how this is grounded in logic, so much as, perhaps, intuition.

---

[37] Although it is not clear from the relevant passages that Frege's remarks are intended to provide necessary and sufficient conditions, Burge (2000, 359–61) convincingly argues that they are. I think it matters not much for present purposes, though see Bar-Elli (2010) for a dissenting voice.

To show how arithmetic is reducible to logic, Frege adopts an axiomatic method (subsequently to be very influential), wherein certain general, basic logical laws are the axioms, and other rules, definitions, and ultimately arithmetical equations can be derived from them, as theorems. Take "7+5=12". The axiomatic method requires us first to define the components of the equation (the numbers and the operator), then offer the proof of the equation—its "ultimate justification"—the premises of which, for the equation to reduce to logic, must be logical truths (1884, §3). In the fullest extension of logicism, then, this approach is generalised: all arithmetical concepts (such as "number") need to be defined in terms of logical concepts (such as "implication", "negation", and so on). The basic logical laws taken as premises in these deductions are thus the "axioms" of Frege's system.

For Frege, these axioms, the so-called basic truths—general logical laws—"neither need nor admit of proof" (Frege 1884, §3), because he takes them to be self-evident. As Kanterian has argued, though in the *Begriffsschrift* Frege seems wary of claiming that the general logical laws are known through some intuitive faculty, by the time of the *Foundations* (1884) it looks as though Frege is opening up to the idea (Kanterian 2012, 25ff).[38] But the type of intuition involved in immediately recognising logical laws to be true must not, for Frege, be the sort of intuition involved in recognising general geometrical truths—"spatial intuition" (Kanterian 2012, 40f)— since, as stated above, the logicist project is epistemological. Accordingly, the aim is to show that arithmetic is *justified* by being based on non-intuitive content; this is why the self-evident character of logical laws, Frege's axioms, are described by him as maximally reliable (Kanterian 2012, 26), unlike flimsy intuitions.

So Frege's notion of analyticity, I have said, is best understood within his logicist project, and as discussed this project proceeds by an axiomatic method to show that arithmetic derives from logic. Analyticity is connected with this system in the following way: when a particular truth can be shown to rest only on logical laws, rules of inference derived from them, and definitions, then for Frege it is analytic. Here is Frege's definition of the analytic-synthetic distinction:

> When a proposition is called . . . analytic in my sense, . . . it is a judgement about the ultimate ground upon which rests the justification for holding it to be true. . . . The problem becomes . . . that of finding the proof of the proposition, and of following it up right back to the primitive truths. If, in carrying out this process, we come only on general logical laws and on definitions, then the truth is an analytic one, bearing in mind that we must take account also of all propositions upon which the admissibility of any

---

[38] One wonders also whether a logical intuition for Frege has a parallel in Russell's idea of acquaintance as applied to logic: "in thought we have acquaintance with objects of a more abstract logical character" (B. Russell 1905, 479).

of the definitions depends. If, however, it is impossible to give the proof without making use of truths which are not of a general logical nature, but belong to the sphere of some special science, then the proposition is a synthetic one. (Frege 1884, §3)

The terms "analytic" and "synthetic" thus apply, for Frege, not to the content of a true proposition, but to the manner of a proposition's *justification*—i.e., propositions are either analytic or synthetic according as they have the different justifications. Take "7+5=12". Where Kant held that since the concept of "12" is not contained in the concepts "7", "5", or their union, the proposition (or judgement) is synthetic, Frege holds that because its components, once defined, are deducible only from general logical laws (such as the law of identity or the law that "if *x*, then whatever else *y*, *x* follows"), the proposition is analytic.

There is an important difference, then, between Frege's and Kant's definitions of analyticity, since for Kant, the analytic-synthetic distinction properly applies to the content of judgements. That is shown by the fact that Kant specifies a logical criterion of analyticity, which concerns the containment-relations between concepts (though it is true that analytic judgements have, on Kant's interpretation, an epistemological difference to the synthetic judgements, due to their containment-relations). Moreover, Kant is not wholly concerned with a circumscribed area of knowledge, such as arithmetic, and so analyticity is not relative to such a reductive system of proof. Frege's placing analyticity in the context of logicist systems will, however, ultimately change how many work with the notion after him.

In any case, given Frege's definition of analyticity, one may wonder whether Frege's axioms, the general logical laws, are analytic. But, as Kanterian points out, since for Frege only a truth derivable (and thus knowable) from general logical laws and definitions can be analytic, and since an a priori truth is one derived exclusively from general laws, for Frege logic itself is neither analytic nor a priori (2012, 24). However, Dummett argues it is essentially an oversight on Frege's part that his definition of analyticity did not extend to the basic axioms of his system (1991, 24). And contra Dummett, Bar-Elli suggests that though these basic truths do need to come out as analytic for Frege's logicist project to get off the ground, and though Frege indeed admits that general logical laws are in no need of proof, Frege must have conceived another method of justification besides proof, such that the basic axioms definitely do indeed come out as analytic (Bar-Elli 2010).

In my view, it is not clear at all why Frege's axioms need to be analytic: his logicism was an epistemological, reductionist project that sought to provide the firmest foundation for arithmetic; the axioms' self-evident character would be thought to provide the sought-after epistemological certainty. Self-evidence is thus a constraint on the aptness of a true proposition for being an axiom. (See Joan Weiner (2007, 681–82) on other likely constraints we can infer from Frege's writings;

see also Burge's (2003a) counter to Dummett that Frege's characterisation of basic logical laws is not an oversight but deliberate.)

Beyond Frege's thinking of his basic logical laws as self-evident and neither a priori nor analytic, there is one further characterisation of them which proved influential, and is of critical importance for this thesis. Namely, he conceived the laws as *normative*. In the introduction to his *Basic Laws*, Frege laments that this point is too easily lost by the psychologistic logicians of his day:

> The ambiguity of the word "law" here is fatal. In one sense it says what is, in the other it prescribes what ought to be. Only in the latter sense can the logical laws be called laws of thought, in so far as legislate how one ought to think. Every law stating what is the case can be conceived as prescriptive, one should think in accordance with it, and in that sense it is accordingly a law of thought. (Frege 1893, 1:xv)

The temptation, in Frege's eyes, has been to construe "law" in its descriptive sense, and then to think that logical laws are laws related only to individual humans. This opens up the possibility that the laws could be otherwise, and fails to capture the normativity of logical thinking: if we want our thoughts to *be* true, then we must follow the most general basic logical laws.

Frege here signals his commitment to a universal, unchanging logic. In the decades to come, others, such as Russell, Wittgenstein, and early Carnap would share in this universalist conception of logic, yet later some logicians would forgo this view (see my §2.2.3). Even so, the idea that Frege explicitly links to the universalist conception—logic's being prescriptive, normative—would remain for some time.

§2.2 Logical empiricism

Unfortunately for Frege, his logicist project came to a premature end, after the now infamous blow dealt by Russell's paradox. Without going into detail (since the fate of his logicism is not our main concern), the paradox made one of Frege's "Basic Laws" untenable, forcing him to stall the project. Russell, however, took up the torch, along with Whitehead in their *Principia Mathematica* (1910). This huge work set about grounding mathematics in logic without succumbing to the paradox, by means of a theory of "types", which is not our concern to explore here.

Whatever the status of Frege's and Russell's logicism, one thing that remained, and indeed bloomed, in philosophical thought over the following decades was a commitment to the use of the new logical calculus. Though Ludwig Wittgenstein parted with Russell and Frege on the idea that logical laws are self-evident, his first major work, the *Tractatus Logico-Philosophicus*, signalled commitment both to the new logic and an attendant decompositional and transformative method

of analysis, partly inspired by Russell's logical analyses. "Self-evidence," he had written, can be "dispensable in logic, only because language itself prevents every logical mistake.—What makes logic a priori is the *impossibility* of illogical thought" (1921, §5.4731; cf. Hacker 1986, 44–45 for discussion). Like Frege, then, though for different reasons, Wittgenstein believed in the universality of logic. Moreover, Wittgenstein brought fresh insights into the nature of logical truth: since logical propositions are true under all conceivable circumstances, they are, he held, tautologous and empty of factual content, and this explains their necessity. Consequently, as Hacker (1986, 46) points out, Wittgenstein thought we do not "know" logical truths, since they had no sense (which is not to say that they are *un*known or nonsensical). Moreover, Wittgenstein thought that the equally tautologous status of all logical truths undermines the axiomatisation of logic: none are especially primitive, for instance because self-evident as Frege's (and Russell's) axioms were regarded (Hacker 1996, 33–34; Wittgenstein 1921, §6.127).

Wittgenstein's idea that logical truths are tautologous was to prove extremely influential; the use of that idea to quash logicism, however, less so. Specifically, the tautology—the emptiness—of logical truth, along with other aspects of the *Tractatus*, became pivotal to the project of a group of Viennese philosophers[39] (and other philosophically minded thinkers), who dubbed themselves the "Vienna Circle". The Circle ambitiously attempted to use the new logic to ground the entire edifice of human knowledge, now emboldened by the possibility of extending Wittgenstein's view of logic as contentless to the whole of mathematics (Awodey and Carus 2007, 26). Accordingly, rather than following Wittgenstein's negative lead on the matter, they instead beckoned in a new wave of logicism. The philosophy they espoused came to be called "logical positivism"; eventually, "logical empiricism".

§2.2.1 The Vienna Circle

Logical empiricism was committed to a few central ideas: to the use of the new logic, a radical empiricism, a critique of metaphysics, and the so-called "verifiability criterion of meaning" (see below). Further, they also shared a broad view of analytic truths as expressing linguistic conventions, especially as applied to necessary and putatively metaphysical propositions. Circle member Moritz Schlick's *General Theory of Knowledge* (1925; German-language first edition published in 1918) had prefigured this focus on convention. Schlick took inspiration, in turn, from Hilbert, who argued that the discovery of new non-Euclidian geometries showed that the axioms of

---

[39] In particular Carnap, who regarded the tautology of logical truth as an especially important insight he learned from Wittgenstein; see Carnap's "Intellectual Autobiography" (1963, 25).

geometry are chosen (arbitrarily, so long as consistent), and thereby that they "implicitly define" the constituent expressions within the axioms (such as "line", e.g.) (Hacker 1996, 46).

The new logical calculus seemed to be the perfect tool to fulfil Enlightenment ideals of articulating what knowledge really consists in, without succumbing to subjective, Romanticist notions. Thus, the Vienna Circle believed a fully deductive approach using the new logic could reconstruct our folk concepts from the ground up, thereby issuing scientifically respectable alternatives in their place (Carus 2007, 14). One tool of this reconstruction task was the aforementioned verifiability criterion of meaning. The Anglophone representative of logical positivism,[40] A. J. Ayer, expounds the principle in the preface to his *Language, Truth and Logic*:

> To test whether a sentence expresses a genuine empirical hypothesis, I adopt what may be called a modified verification principle. For I require of an empirical hypothesis, not indeed that it should be conclusively verifiable, but that some possible sense-experience should be relevant to the determination of its truth or falsehood. If a putative proposition fails to satisfy this principle, and is not a tautology, then I hold that it is metaphysical, and that, being metaphysical, it is neither true nor false but literally senseless. (Ayer 1936, 9)

(The final sentence here is not itself part of the principle, but a hypothetical application of it.) We should note here that it is a matter of controversy which kinds of proposition—or indeed, whether propositions or sentences—play the verifying role: those reporting sense-data, as Ayer seems to believe, or those describing physical objects.

Problems aside, the basic idea is that when we find ourselves unable to locate the possible sense-experiences (or whatever else) which would determine whether or not some proposition, such as "God is omnipotent", is true, and having also found it not to be tautologous, we must declare it as void of meaning. Only propositions which are tautologous or verifiable, on this criterion, are meaningful. Whereas, in contrast to religious, spiritual, aesthetic, and moral propositions, empirical propositions are easily conceived to be verifiable by some (set of) sense experiences. It will help to have an example. With a little consideration of the statement "Tomorrow at noon, the temperature in the first floor of the university library will be 25°c", we can readily imagine the sorts of sense experiences (or material object statements) with which to test, in different ways, its truth or falsehood. We know that a measuring device of a certain sort placed in that location (specified more fully than in this proposition) at that given time would, if the statement is true, show the measure reaching a mark indicating the number 25; if false, it would

---

[40] Recent Carnapian scholarship does not look too fondly on Ayer, especially insofar as his interpretations of the Circle's philosophy are concerned. See Friedman (1999) and Carus (2007).

read another number. Further, as part of the reconstructive program, apparently subjective terms such as "hot" and "cold" could be given, it was thought, precise reconstructions (Carus 2007, 15).

Analyticity enters the picture here, since, for the Circle, empirical truths are those made true by particular experiences, while logical truths, being tautologous, had to be true by virtue of their form alone: "made" true by the conventions of the symbolism (see Hacker 1996, 47). Further, as we discuss in §3, one way to try and define analyticity is via synonymy, and the verifiability principle of meaning also offers a criterion for determining synonymy: since we can give the meaning of an empirical statement by specifying the conditions under which it can be made true, statements have the same meaning (are synonymous) insofar as they share the same verifiability or falsifiability conditions.[41]

The above discussion sketches out the general commitments of the Vienna Circle, but in this section henceforth, I want to focus on Rudolf Carnap. For one, analyticity remained pivotal to his changing philosophical projects; for another, his debates with Quine formed the foundation for much of what became analytic philosophy in the decades after (cf. Hacker 1996, 187-89). I will outline in brief each of Carnap's projects, and show how each corresponding analytic-synthetic distinction adapted to those projects.

§2.2.2 Carnap: The *Aufbau*

Carnap's *The Logical Structure of the World* (1928)[42] was a bold attempt to carry out the logicist ambitions of the Circle by securing the foundations of mathematics and the sciences. The system he advanced in the *Aufbau* attempts to reduce the language of physics down to a basic level:

> [I]t attempts a step-by-step derivation or "construction" of all concepts from certain fundamental concepts, so that a genealogy of concepts results in which each one has its definite place. It is the main thesis of construction theory that all concepts can in this way be derived from a few fundamental concepts. (Carnap 1928, §1)

Accordingly, so-called construction theory aimed to reduce mathematical, logical, and scientific statements down to statements which no longer contained the original, problematic concepts, showing them to be constructed from the fundamental concepts. The reductions of the *Aufbau*, then, were somewhat like translations, recalling the transformative model of analysis discussed in §1. However, these apparent translations are such that one language (the formal language) is to be

---

[41] This is, at least, Quine's understanding (1951, 37).

[42] German title: *Der logische Aufbau der Welt*; hereafter, simply "*Aufbau*".

more fundamental than all others, which brings to mind the decompositional method of analysis. Both methods of analysis seem to be present in Frege's own logicist project also. (Strictly speaking, Carnap's translations or analyses were *quasi*-analyses, though the distinction is not important for our purposes; see Richardson (1998, 51–59) for discussion). Further, the verifiability criterion of meaning finds application within this reductive system: any concepts which cannot be reconstructed from the base language can duly be proclaimed "nonrational" and "metaphysical" (1928, §176).

The nature of the *Aufbau* reductions is contested (though not that the *Aufbau* is reductivist). An older interpretation sees Carnap as a radical empiricist who wanted to construct objects out of only spatiotemporal sense-data points, conceived as epistemologically privileged—the first, justified steps of knowledge. As Juhl and Loomis point out (2009, 24–25), there is some evidence that Carnap himself interpreted his *Aufbau* project this way. The newer interpretation, however, represented by Michael Friedman (1987, 1999), A. W. Carus (2007), and Alan Richardson (1998),[43] among others, maintains that this picture is misguided for reasons we explore presently.

For one, Carnap does not take sense-data as given or primitive; instead, he builds up *to* them after a very lengthy construction which begins (and mostly stays) in the autopsychological domain (Friedman 1987, 522–23). The term "autopsychological" denotes, as Carnap describes it, "my" experiences; these do include, indeed, such experiences as a shade of blue in a particular place in the visual field. However, Carnap is keen to point out that these autopsychological terms are not meaningful in advance of the system's construction—in advance, that is, of the nonpsychological, such as the physical, and the "you" having also been constructed (Carnap 1928, §65).

For another, Carnap is open in the *Aufbau* about the legitimate bases for constructional systems: they could just as well be physical as autopsychological, and each have different advantages. Foreshadowing aspects of his later work, it is a matter of choice which base we build from (Carnap 1928, §62; Friedman 1987, 524; Richardson 1998, 24). (This aspect of the *Aufbau* is fairly open, then, although the logical side is not: the formal language to be used in the constructions is conceived as a universal logic, that of Russell and Whitehead's *Principia*; this will soon change for Carnap.) Furthermore, Richardson (1998, 23) argues that the early project is guided by the epistemological quest for objectivity, having accepted that the logical constructions will begin from subjectivity—the "given" in experience—which, contra traditional empiricism, does not qualify as knowledge.

---

[43] Richardson (1998, 2–3) provides a good list of sources for the late-twentieth century shift in approaches to the work of the logical empiricists in general.

Whatever the best way to interpret and characterise Carnap's early project, his constitutional definitions required much intellectual craft and finesse: the *Aufbau* painstakingly attempts to construct, step-by-step, a description of the world of physics from slender beginnings—from, at base, sentences about the recollection of similarity between elementary experiences, and via, *inter alia*, definitions of the individual sense modalities and colour qualities.

Analyticity features in this grand project when characterising the theorems of its constructional (or constitutional) systems; the theorems, Carnap explains, split into analytic and empirical. Any such system reconstructs a concept from a formal base; this process thus begins with a definition or statement of the construction in the "base logistic language". The formal definition is subsequently followed by re-statements in three auxiliary languages: the "paraphrase", "realistic", and "fictitious operation" languages (1928, §106). Carus notes that these auxiliary languages essentially ease our transition into the formal language, whose statements we cannot accurately translate at the informal level:

> The new, reconstructed concepts could still be described or explained or gestured at in ordinary language, or even in traditional philosophical language. But such forms of discourse can only be a more or less inaccurate approximation, a user interface for human users of genuine knowledge, whose precise and canonical statement is in the standard logical language. (Carus 2007, 14)

Thus, natural languages such as German and English were conceived at this stage as imprecise tools, a means of accessing the formal language. This was nothing new by Carnap's time, of course. Much of this viewpoint was common to Frege:

> In the first stages of any discipline we cannot avoid the use of ordinary words. But these words are, for the most part, not really appropriate for scientific purposes, because they are not precise enough and fluctuate in their use. (Frege 1914, 207)

(Indeed, the similarities do not end there: in the same text, from a lecture, Frege discusses "constructive definition" (1914, 210)—which allows one to introduce a new sign to take over the complex sense of an older sign—for purposes of constructing a system "from the ground up".)

Anyway, with Carnap's idea of a reconstruction in view, we can now approach his *Aufbau* definition of the analytic-synthetic distinction:

> The statements or theorems of a constructional system are divided into two different types . . . The first type of theorem can be deduced from the definitions alone (presupposing the axioms of logic, without which no deduction is possible at all). These we call *analytic* theorems. The second type of theorem, on the other hand, indicates the

relations between constructed objects which can be ascertained only through experience. We call them *empirical* theorems. (Carnap 1928, §106)

Shortly, we will absorb this definition by considering its application in practice to some of Carnap's actual constructions. At present, it is apposite to consider in brief its relation to Kant. Just after offering this definition of analyticity, Carnap refers to Kant's own distinction, chiefly to point out that he rejects the synthetic a priori; there are "no such" judgments, according to construction theory. The matter is perhaps less clear, however, given the function of these judgements, as Richardson discusses: "The epistemological role that Carnap assigns to the logical truths embodied in his constitutional definitions is the methodological role played in Kantian philosophy by synthetic *a priori* principles: They first make possible objective, empirical knowledge" (Richardson 1998, 196). At any rate, given Carnap's above definition, he has no room for the synthetic a priori; it would be an invalid type of theorem or statement.

In §108, Carnap formalises the given at the base of his constructions, which is a basic relation denoted by "Rs" in the base logistic language;[44] in the paraphrase language, he calls this the recollection of similarity between elementary experiences. From this basic relation, Carnap can then derive theorems. The first theorem Carnap states is "Rs ∈ as", which means (in the paraphrase language) that the recollection of similarity (Rs) relation between two elementary experiences is asymmetrical (which is symbolised as "as"). To understand what this means, consider the following two elementary experiences: "*y*", that of experiencing an amber colour patch before one, and "*x*", that of experiencing the memory of an amber patch. By "asymmetrical", then, Carnap means that *x* may be a recollected experience of—may be partly similar to—*y*, but not the other way around, since there is a temporal order to them.

Now, we can see the theorem "Rs ∈ as" (that the recollection of similarity between experiences is asymmetrical) does not follow from any definition. Accordingly, Carnap classifies it as empirical (or synthetic). By contrast, the construction "part-similarity" (§109) in turn is defined[45] as follows: "if a recollection of similarity holds between elementary experiences *x* and *y*, then a part of *x* is similar to a part of *y* and a part of *y* is similar to a part of *x*". It follows from this definition (when written in the base logistic language) that "part-similarity" is symmetrical; i.e., when one experience is recollected to be similar to a memory experience, those experiences are to that extent (at least)

---

[44] Carnap intends the *Aufbau* to articulate the general constitutional system of the whole of science. Thus, this relation is to be the very base of the constitutional system of scientific knowledge (Richardson 1998, 8–9).

[45] Strictly speaking, the definition is always in the base logistic language.

part-similar. Likewise, every experience is part-similar to itself. Accordingly, the derived theorems "Ps ∈ sym" (meaning that the relation of part-similarity is symmetrical) and "Ps ∈ refl" (meaning that the relation is reflexive) are *analytic*: they are deducible from the construction's definition alone. It is thus that an analytic-synthetic distinction finds a role in Carnap's *Aufbau* project.

Before moving on, two points are worth noting about Carnap's characterisation of analytic and synthetic statements. First, he describes analytic theorems as "tautological statements about concepts", while "the empirical theorems express an empirically ascertained state of affairs" (1928, §107). So it seems that at this stage, Carnap echoed the Circle's guiding thought that analytic truths are tautologies, and further that "*[l]ogic (including mathematics) consists solely of conventions* concerning the use of symbols, *and of tautologies* on the basis of these conventions" (1928, §106). Second, whether some theorem counts as analytic or not will, for Carnap, depend on the how the construction is defined at the base level, which introduces an element of choice.

Taking these two aspects into account, some comparisons with Frege's conception are forthcoming. First and foremost, Carnap characterises analyticity by a criterion which certainly recalls Frege's own approach—perhaps this is not so surprising, given that both definitions feature in logicist projects. Further, as we saw in §2.1, a statement is analytic for Frege when it can be derived from logical laws and definitions alone; Carnap notes that we "presuppose" the axioms of logic, and so his conception of analyticity appears to be very much influenced by Frege at this stage in this thinking. Likewise, on Carnap's approach, analytic theorems seem relative to the chosen definitions at the base of a construction, which in principle could vary, or so it seems to me, though both Frege and this early Carnap are committed to the logic of the *Begriffsschrift*.

§2.2.3 Carnap: *Syntax*

By the time of *The Logical Syntax of Language* (1934), Carnap's views had changed in some important respects. Significantly, he became aware of Kurt Gödel's incompleteness theorems, the first of which was published in 1931. Gödel had shown that for any axiomatic system, such as a logicist construction of arithmetic, there will always exist a true sentence formalisable in that system (perhaps, say, Goldbach's Conjecture that every number greater than two is expressible as the sum of two primes) which we cannot *prove* to be true (or false) within the system; i.e. we cannot derive the statement from the axioms. We could always prove (or disprove) the statement trivially, by making it (or its negation) an axiom of the system, but then Gödel's proof taught us that yet another statement will always exist which is not provable within this newly expanded system. Thus he had shown that no consistent system could be complete; moreover, that no complete system could actually demonstrate its own consistency—i.e., could show it was without contradiction.

Gödel's proofs changed things drastically for Carnap, as well as for logicism. As Steve Awodey notes, the results immediately cast into doubt the very idea of explaining all a priori knowledge in terms of analyticity or logical truth, insofar as "analytic" means tautological or trivial (Awodey 2007, 226–27). Perhaps that alone may not be a problem for, say, Frege, since (as we discussed) he did not question the self-evident status of his axioms. Yet more generally there is a problem for logicism here, since in both the *Aufbau* and Frege's *Begriffsschrift*, a statement's analyticity is characterised by its demonstrability from (at least) axioms, and Gödel shows that arithmetical truths will always exist which are not demonstrable in the basic logic. Accordingly, since Gödel's incompleteness results had shown there was no way forward for a logicism trying to *prove* all true statements of mathematics within a finite system, Carnap sought a wider notion of logical truth than is allowed by provability. In essence, then, he looked for a new definition of analyticity.

We can see how a new definition struck Carnap as possible by means of a useful analogy.[46] On a chessboard, pieces are arranged in starting positions, which we can imagine as the axioms of a formal system; further, chess rules govern permissible moves, which we can analogise as inference rules; and finally, series of moves lead to positions of checkmate, and we can think of these entire move sequences as theorems. The idea that there could be some stronger conception of logical truth than provability, such as analyticity, is here the idea that there are configurations of pieces constituting checkmate which cannot be achieved through a permitted chain of moves. One such configuration is a row of black pawns, each unmoved from their starting positions, and, behind them, a white rook in one corner checking a king in the other. Thereby, we have a position which looks to be a legitimate checkmate (a true statement of the system), but which we cannot ever demonstrate as culminating via a chain of valid sequential moves from the starting positions (i.e., as provable within the system). Nonetheless, the configuration remains checkmate, like a statement that is necessarily true but not provable from the axioms.

What the chess analogy shows is that a new definition of analyticity would have to be formed syntactically—i.e., according to the rules of a formal system (not referring at all to the meaning or interpretation of its signs), much as the checkmate is valid by virtue of the rules of chess. Carnap soon found out, however, that any attempt to define analyticity syntactically within an axiomatic system would encounter fatal problems (Awodey 2007, 228), and so Carnap followed Tarski, Hilbert, and Gödel in adopting "metalinguistic" systems, able to talk about the original "object" language in a so-called meta-language. To avoid contradictions, Carnap had to employ a meta-language more resourceful or "stronger" than the object-language (the language of science); within

---

[46] I take this illustrative analogy directly from Awodey and Carus (2007, 36).

this more expressive language, then, he could define analyticity for the scientific language in terms of its rules. And so, in the *Syntax*, he finally found a way to characterise analytic or logical truth which was not vulnerable to incompleteness worries because not cast in terms of provability within a system.

Specifically, in the *Syntax*, Carnap now uses the meta-language to define analytic sentences for the object-language as the consequences of every sentence (or as the consequences of the "empty set" of sentences, which is equivalent), thereby picking out those sentences which are true whatever the empirical facts are (1934, 39). Thus, analytic (or logically true/"L-true") sentences are not necessarily demonstrable from the axioms, and are instead classifiable as "L-determinate". (Likewise, contradictory sentences, which have every sentence as a consequence, are L-determinate; a sentence is thus determinate if and only if it is either analytic or contradictory.) Carnap then uses determinacy to characterise logic more generally (Carnap 1934, 39–42, 177); this manoeuvre allows him methodically to separate out logical from empirical truths while no longer being shackled by concerns about provability within a system.

However, having now ascended to the meta-linguistic level, Carnap no longer saw it necessary to commit himself to the universalist conception of logic he took over from the *Tractatus*. If logical truth can only be defined *for* an object-language, such as the language of science, then the question over the "right" definition of logical truth seems to be ill-formed. Of course, he realised in turn, different logical systems must be possible, since logical truth or analyticity is always relative to a language system. This realisation became the famous "principle of tolerance". (I think it is fair to say, however, that there were shades of the tolerant approach in the *Aufbau*, especially if the more recent interpretation of the *Aufbau* prevails.) Carnap expounds the principle in the *Syntax*:[47] "It is not our business to set up prohibitions, but to arrive at conventions" (1934, §17). He continues:

> *In logic, there are no morals.* Everyone is at liberty to build up his own logic, i.e. his own form of language, as he wishes. All that is required of him is that, if he wishes to discuss it, he must state his methods clearly, and give syntactical rules instead of philosophical arguments. (Carnap 1934, §17)

The principle of tolerance is thus Carnap's way of freeing himself from a universal logic. Which logic one "chooses"—say, intuitionist or classical logic, will be decided by pragmatic criteria as to which is more fruitful. Accordingly, for Carnap by this stage, no single logical framework governs rational thought (Friedman 1999, 169).

---

[47] Richard Creath (2009) argues the principle itself came relatively late into the writing of the *Syntax*, however.

The impact of tolerance on analyticity is significant. While, in the *Aufbau*, sentences are only analytic relative to the formal framework in which they occur, now, with the principle of tolerance in play, many different frameworks are possible, such that quite different sentences could be analytic given the different available frameworks. Likewise, it is only within a logico-linguistic framework that synthetic sentences are possible, since the analytic-synthetic distinction is provided by the framework's logical structure. Consequently, as Richardson argues, it is not quite right (as we may be tempted) to say that logic, for Carnap, is fact-independent: every language has its own logical structure, and thereby its own analytic-synthetic distinction which delimits what, ultimately, could even count as a fact in the first place (1998, 217).[48]

Finally, this pluralism thus extends to working out how best to structure the language of physics, which is still Carnap's aim. That is, the kind of sentence sitting at the bottom of a reduction, a so-called "protocol sentence"—say "there is green before me now"—is open to decision, constrained only by practical criteria (such as relative syntactical complexity) and not by philosophical concerns as to ultimate correctness (Carnap 1932, 465). But although the construction of scientific language remains the main aim, Carnap no longer conceives the project in epistemological terms: he is not working out how to move from a subjective, autopsychological language to objective (or intersubjective) scientific language, as he was in the *Aufbau*. Rather, the protocol sentences are already objective, and autopsychological language is not epistemically privileged (Richardson 1998, 216).

Further detailed discussion of Carnap's (rather technical) *Syntax* project, and the employment of analyticity within it, goes far beyond the scope of this chapter. We can, however, briefly mention some interesting changes. First, since Carnap disowns the epistemological commitment of the *Aufbau*, analytic sentences are no longer those which are derived from a subjective base language (out of which objective language must be constructed). Second, given the new pluralistic approach, logical empiricists may choose the (base) protocol language as they like, abiding only by pragmatic criteria; this means that the analytic-synthetic distinction is now relative to the chosen formal object-language. A further consequence of this fact is that the analytic-synthetic distinction itself cannot be defined by some epistemological criterion: epistemic notions, such as confirmation, are defined within a framework, and the framework's very principles are expressed by analytic sentences. Thus, it is inaccurate to say that for Carnap analytic sentences are those which have some special epistemic authority in the sense of epistemological access (Richardson 1998, 223);

---

[48] Indeed, this feature of Carnap's notion of analyticity mirrors the role of the synthetic a priori for Kant; as Richardson points out, Carnap's analytic sentences condition the possibility of empirical judgment, all relative to a given language (Richardson 1998, 227).

the only priority they enjoy is logical, not epistemic—this point will have significance in both §3 of this chapter, and the next chapter of this thesis.

§2.2.4 Carnap: Semantics

In Carnap's later period, analyticity still remained essential to his thinking. This period was characterised by a turn toward semantics, and eventually the notion of explication. Heavily influenced by Tarski and his disquotational truth-schemas (of the form " 'S' is true iff p"), Carnap employed Tarski's semantical methods to define in the meta-language (for the formal object-languages) such notions as "truth" and "designation"—notions he disallowed in the *Syntax*. Over time, Carnap became interested in employing semantical methods in order to "explicate" troublesome concepts, such as those of truth, designation, and analyticity itself. Carnap explains in *Meaning and Necessity* (1947) that in an explication he replaces a familiar though vague concept (the "explicandum") with a precise alternative (the "explicatum").

The explicatum of the apparently vague concept of analyticity is "L-true", which, he notes, "is meant as an explicatum for what Leibniz called necessary truth and Kant analytic truth" (Carnap 1947, §2). Presumably, given our exposition in §1, Carnap has in mind what may best be described for Leibniz as analytic-necessary truth, since (if Leibniz does have a comparable conception of analyticity), ordinary factual truths are also analytic for God, in Leibniz's view, although contingent. Analyticity for Kant is a comparatively simple affair, though I think since Carnap would find conceptual containment far too ambiguous, he likely has in mind Kant's epistemological and logical characterisations; respectively, that an analytic judgement must not be ampliative and that its denial is a contradiction in terms.

The relationship between explication of analyticity and L-truth is notable. By this later stage in his thinking, Carnap clearly still regards analyticity as an important concept, but he deems it in need of transformative clarification. Yet for there to be an explicatum of analyticity, there must be, as Carnap accepted, an everyday pre-theoretical explicandum—i.e., not one of his prior definitions for a formal language, but one in English, say, or German. Beyond those, however, he does appeal to the theoretical conceptions as formulated by Kant and Leibniz; perhaps Carnap regards these as explicata of more ordinary notions, though not explicata which were exact enough. "Enough" is to be cashed out in terms of practical purposes—fruitfulness (as well as exactness) is one of the desiderata which constrain the process of explication, for Carnap.[49] Another is similarity,

---

[49] Carnap discusses the four desiderata of explication in §1 of the subsequent work, *Logical Foundations of Probability* (1950).

so presumably Carnap had seen in ordinary talk, in philosophical history, and in his own past formulations some kernel of analyticity which he wanted to reproduce, but put to a more fruitful purpose, given his motivations. In any case, it is only due to this "given" explicandum of analyticity that Carnap can provide a definition for its explicatum, L-truth. The definition is as follows:

> 2-2. *Definition.* A sentence S_i is *L-true* (in S_1) $=_{Df}$ S_i holds in every state-description (in S_1). (Carnap 1947, §2)[50]

Here, the notion of a "state-description" eliminates and replaces talk about logical truth or analyticity. A state-description is a class of sentences in an object-language (here, the language "S_1") which contains (and only contains) every atomic sentence of that language or its negation (Carnap compares state-descriptions to Leibniz's possible worlds and early Wittgenstein's "states of affairs"; essentially, a state-description is a line in a truth-table). A state-description thus describes a possible state of the respective universe of discourse, and if a sentence holds in a state-description, this simply means that were the world arranged that way, the sentence would be true. Now, grouping every state-description together conjointly would describe all the possible ways the world could be arranged. For example, taking classical logic as our framework language, Arnold could be proud ("Pa") or he could not be proud ("¬Pa"); in any given state-description, either of those two sentences will be true (within classical logic, at least). Accordingly, "Pa v ¬Pa" holds in every state-description, and is thus L-true. And since such a sentence is true in every state-description, we have no recourse to the "facts of the universe" in order to demonstrate its truth. This means in turn that Carnap's definition of L-truth satisfies his general motivating idea of analyticity, which he spells out in an adequacy condition:

> 2-1. *Convention.* A sentence S_i is *L-true* in a semantical system S_1 if and only if S_i is true in S_1 in such a way that its truth can be established on the basis of the semantical rules of the system S_1 alone, without any reference to (extra-linguistic) facts. (Carnap 1947, §2)

This convention is plausibly that very kernel of prior formulations which Carnap wanted to retain in his explications: the basic thought that some truths, relative to a logico-linguistic system, are true because of the rules of that system, and thus make no mention of facts. This motivating idea is that which Quine will object to the most, and the underlying distinction between facts and logic that it relies on. Indeed, we turn to Quine's criticisms shortly.

---

[50] The term "S_i" is a variable depicting an object-language sentence; the term "S_1" denotes one of Carnap's symbolic language systems (1947, §1).

Having surveyed the different projects to which Carnap was committed, we can see that analyticity was pivotal to each of them. Further, the main differences between the conceptions are those which allow ways for certain core ideas to persist, given new obstacles. For example, throughout these projects, analyticity always remains relative to a language system, but that relativity is cashed out differently in response to challenges and opportunities presented by contemporary thought at the time. Likewise, Carnap always wants to isolate the notion of logical truth for the purposes of grounding the language of science. Both commitments remain central through successive stages of his philosophical career.

In the *Aufbau*, a sentence in the base logistic language is analytic when it follows from the axioms of logic and arithmetic; true, the only logical language he has in mind at this stage is classical, but the analytic-synthetic distinction is still a technical device relative to the framework we employ, and it allows us to tell apart the theorems whose truth we can verify by empirical means from those we cannot. However, Gödel had shown that we cannot derive from such a formal system's axioms every true statement, and so Carnap adapted analyticity in response. In the *Syntax*, no longer wedded to the universality of classical logic, Carnap argues that sentences are analytic relative to any of the different logical languages we choose, so long as their rules are clearly and exactly spelt out. Carnap thus employs the analytic-synthetic distinction only when using a resourceful meta-language for a given object-language: an object-language sentence can then be specified as analytic when it follows from every one of its sentences; i.e., when it is a consequence of the empty set. Thereby, the thought once again that analyticity and logical truth precede empirical matters is retained.[51] Finally, that crucial separation is just as relevant in Carnap's later period, too, where he now uses a meta-language to specify a sentence as L-true in the object-language iff it holds in every state-description of that language. Accordingly, there are no possibilities in which that sentence is false, such that factual matters are not relevant to its truth.

There is much left to say about the success of Carnap's use of the analytic-synthetic distinction as it pertains to the goal of his projects. We cannot engage with most of the appraisal even from the literature, but we would be remiss not to recite and discuss Quine's basic objections to Carnap's use of the distinction.

---

[51] I choose "precedes" rather than "do not depend on" to be mindful of Richardson's point as discussed, in my §2.2.3, that discussion of factual matters is only intelligible within a pre-existing logical framework, for Carnap.

§3 Quine

The analytic-synthetic distinction enjoyed a heyday under the logical empiricists, and especially under Carnap. But it fell into disrepute not long after, as is well known, due to the criticisms of Quine. It is apt to reflect, briefly, on Quine's legacy in this respect.

§3.1 The impact of Quine

Although the distinction is still employed in contemporary analytic philosophy, it is not wisely employed without safeguards. For example, Amie Thomasson (2007) argues that recent intractable metaphysical debates about the composition of objects are dissolvable by consideration of the analytic interrelations between statements about objects and statements about the "simples" of which objects are supposedly composed. But in order to offer that critique, Thomasson first spends a methodological chapter justifying her use of the analytic-synthetic distinction, directly attacking Quine's arguments (2007, chap.2). Such a procedure, after Quine, makes sense, and this really speaks to his influence.

But it seems that more than Quine's arguments themselves, the general feeling which Quine engendered—that the analytic-synthetic distinction could not be taken for granted—is what led to its relative absence in the contemporary setting. For instance, while Timothy Williamson finds he cannot rely on Quine's arguments against the notion of analyticity in his *Philosophy of Philosophy*, because he finds them no longer compelling (2007, 50), he does admit wanting to stay true to something of a Quinean spirit with his own arguments: "There is something robust about 'Two Dogmas of Empiricism': insights remain even when its skepticism towards meaning is stripped away" (2007, 52). Indeed, in the next chapter, I highlight just how a point commonly cited in material about analyticity today has an ancestor in Quine's discussion.

Were Quine's attack to land—were the analytic-synthetic distinction to be disregarded entirely—one would come to see just how important the distinction was. As we have surveyed it in this chapter, the distinction forms a central part of several philosophical projects. Leibniz's metaphysics of substance is articulated by means of the notion of analysis, which may well imply that all truth is analytic, even if not logically necessary. It is highly questionable whether Hume divided truth into analytic and synthetic, yet the distinction he does own, between relations of ideas and matters of fact, helps him formulate in his later philosophy his investigative intentions. For example, if causal sentences clearly expressed relations between ideas, Hume would not find them troubling; yet if, as he suggested, they express propositions concerning matters of fact, then there is ample room for philosophical confusion, which needs clearing up. Hume's distinction

nonetheless influenced Kant, whose conception of transcendental philosophy—its entire point and purpose—is expressed in terms of the possibility of synthetic a priori judgements, which, he held, condition the possibility of experience. Next Frege, as a logician and mathematician, wanted to secure arithmetic in logic—to bring its theorems out from the synthetic cold as Kant had left them; to show arithmetic to be fully analytic. The Fregean project ultimately failed, but Carnap was nevertheless determined to find a way to construct from mathematical and logical axioms the language of science, including physics, arithmetic, and geometry. His various attempts each made use of some conception of analyticity, which unfailingly formed the bedrock of those attempts. There is no doubt, then, of the historical significance of the analytic-synthetic distinction, under its different guises.

Even today, although the term "analytic" is not always happily invoked (with the exception of the "epistemic analyticity" advocates, such as Boghossian, whom we discuss next chapter), there is a sense that some distinction between truth which is conceptual and truth which is empirical must undergird the point and purpose of philosophy as its own discipline. And this is why Williamson's battle against the distinction is the main thrust of his war on what he terms "philosophical exceptionalism" (2007, 3). Contra Williamson, it is an underlying theme of my thesis that there is a compelling way to spell out how philosophy should be thought of as especially conceptual, and so, relatively, exceptional in its turn. But before we can scrutinise Williamson's attack on philosophy as a discipline independent from the sciences, we must finally finish tracing the historical roots of the analytic-synthetic distinction by investigating Quine's influential critique.

§3.2 *Quine's epistemological holism*

Quine's arguments against analyticity invoke his epistemological holism, which he offers in order to criticise the view of confirmation he thinks is implicit in the verification principle of meaning. It is worth noting that this is only one way in which Quine attacks analyticity, albeit perhaps the most direct way, insofar as that notion is put forward by logical empiricists. For this reason, I cover it here before Quine's other criticisms which might threaten not just the analytic-synthetic distinction, but meaning itself.

As elsewhere in his seminal article, "Two Dogmas of Empiricism" (1951), Quine thinks the best route to attacking a given notion is to see what light (or darkness) it casts on related notions. So here he investigates the verification principle of meaning in terms of its effect on synonymy (a notion of which Quine is sceptical to begin with). He thinks that if the verification principle of meaning can offer us a coherent account of when statements are synonymous, it can thereby give us an account of analyticity. Quine infers from the principle that it must have the following relation

to synonymy: two statements are synonymous when (and, I presume, only when) they are confirmed or falsified by the exact same conditions. On this view, then, "Tomorrow at noon, the temperature in the first floor of the university library will be 25°c" means the same as "By midday, it will be 25°c on floor two inside the Templeman building",[52] since the conditions which would verify or falsify the one would equally as well the other.

Yet, as Quine points out (1951, 38), by offering this account of statement synonymy we have tacitly presumed a very particular view of how it is that individual statements, or hypotheses, are verified by experience. Quine sees the paradigm example of this kind of view as embodied in Carnap's *Aufbau*, which, according to Quine, promoted the translation of "significant discourse" into a specified sense-datum language which took statements as primary units of meaning (1951, 39).[53] This explicit reductionism is held by Quine to have engendered a reductivist assumption in the epistemology of his day: "The dogma of reductionism survives in the supposition that each statement, taken in isolation from its fellows, can admit of confirmation or infirmation at all" (1951, 41).

Now, the reason this apparent dogma is relevant is because, Quine claims, it is "intimately connected" with the dogmatic belief in the analytic-synthetic distinction, and in fact that they are "at root identical" (1951, 41). He takes them to be related since he deems the idea of analyticity to depict a limiting case of truth—i.e., given the verification principle of meaning, a statement is meaningful when we can specify the conditions which would verify or falsify it, *except* for those (analytic) statements which are verified "come what may" (1951, 41). Quine demonstrates this limiting case in another way by saying that statements in general are made true by two components: one which specifies, empirically, the way the world is, and the other which specifies the way language is. So, Quine thinks, another definition of "analytic statement" is a statement which has a null factual component. This, as we have seen, does not stray far from the Viennese conception, especially by the point of Carnap's *Meaning and Necessity* definition, that analytic propositions are the consequences of every proposition, which makes them true "come what may" (to use Quine's terms).

At this juncture, I wish to pause and prefigure concerns essential to this thesis (and to the next three chapters in particular). If we commit to this view of "truth-making" in general—that

---

[52] The required context for this similarity to be transparent, of course, is that the phrase "the university library" is a definite description uttered on the campus of The University of Kent, whose library is named the "Templeman building".

[53] There is reason to regard Quine's influential reading of the *Aufbau* as significantly skewed by his awareness of Russell's external world programme of epistemology. The case is made by Richardson (1998).

declarative sentences or other truthbearers are made true by separate linguistic and factual components—then we implicitly assume two deeply problematic thoughts. First, this general view of truth-making tends to presuppose that normative sentences are also descriptive; this is something which I show in the next chapter is present in the work of Chisholm, Boghossian, Williamson, and others (see §2 of my Chapter 2). I find this view very much mistaken, and I take rejecting it to unlock much confusion surrounding the idea of conceptual truth. Second, there is a problem contained in the very idea that a sentence is made true by separable components, language and fact, which beckons a troubling philosophical anxiety about the epistemological relationship between ourselves and the world. Quine certainly succumbs to this second problem, inherent at it is in his picture of intentionality and truth; the relationship between analyticity and this picture will motivate much of Chapter 4. But it is open to question whether we can attribute the problematic first view to Quine, since his epistemological holism, as we see presently, leads him to see the limiting case definition as unintelligible.

Indeed, Quine thinks that this way of framing analyticity is at bottom *nonsense* because it presumes that individual statements are base elements of a larger, divisible whole: "what I am now urging is that even in taking the statement as unit we have drawn our grid too finely. The unit of empirical significance is the whole of science" (1951, 42). Let us use one of Quine's analogies to explain this view. Imagine all humanity's efforts to know the world—be they those in the humanities, sciences, social disciplines, or mathematics—as forming a large fabric. The fabric as constructed will have a peripheral border, and so pieces of knowledge will be scattered about the fabric closer to the border or further back from it. What lies outside the border is, in this analogy, experience itself, so that only some pieces of knowledge or belief actually "touch up" against experience. Where it does this, pieces of belief at the periphery are bound at times to conflict with pieces of experience, which will occasion change:

> A conflict with experience at the periphery occasions readjustments in the interior [of the fabric]. Truth values have to be redistributed over some of our statements. Re-evaluation of some statements entails re-evaluation of others, because of their logical interconnections—the logical laws being in turn simply certain further statements of the system, certain further elements of the field. (Quine 1951, 42)

The idea here is that *any* statement in the system is open to revision, including logical laws (1951, 43). In principle, even statements close to the periphery can be maintained while those we are more committed to near the centre are changed. This possibility, Quine thought, rendered the notion that analytic statements are those that cannot be changed "come what may" completely indefensible.

But, to be fair to Carnap, aspects of this holism and freedom of choice over revision echo some of his own views. For we saw that, at least by the time of the *Syntax*, Carnap thought that the L-true sentences of an object-language were to be specified using a meta-language, and that we had a choice, given the principle of tolerance, to adopt different languages, which thereby changes which sentences count as L-true (analytic in the object-language). Thus, given tolerance, the law of non-contradiction could be true in one language, and not in another. Hylton points out, however, that the principle of tolerance epitomises the problem for Quine: the principle reinforces the epistemological difference between analytic and synthetic truth (it treats some parts of the fabric as entirely separate), since although we are free to choose the analytic principles which in effect constitute the language, our choosing them is supposed to circumscribe them, and prevent them from verification or falsification within the language (Hylton 2016, §3.1). And that, Quine argues, simply is not so.

§3.3 The second dogma

Let us move on to Quine's attack on the other dogma. He begins by covering different ways to draw the analytic-synthetic distinction. Unsatisfied with Kant's containment metaphor, he suggests that it really comes to the following: that a judgement is analytic when true in virtue of meaning and independently of fact (1951, 21). To become more precise, Quine distinguishes between those analytic statements which are logical truths ("No unmarried man is married") and those which are reducible to logical truths by the substitution of synonyms for synonyms ("No bachelor is married"). It is the latter kind on which Quine focuses his critique.

As an aside, note that this latter kind of analyticity, on which Quine focuses, is that which Boghossian more recently termed "Frege-analyticity" (1997, 337). However, in spite of Boghossian's preferred terminology, this is not actually Frege's definition of analyticity, for Boghossian's rendering does not individuate analytic truths according to the method of *proof* involved in their "ultimate" justification.[54] The method of proof requisite to show a judgement to be analytic is one in which we consult only logical laws and definitions. (I think Frege-analyticity is more aptly ascribed to Leibniz, since it does resemble the process as he elaborated on it for demonstrating a truth by analysis; see my §1.1.) Incidentally, Boghossian shows some reservation in a footnote that his so-called Frege-analyticity may not be Frege's, but thinks the important thing "is not who came up with the idea, but rather the philosophical role it has played" (1997, n.13).

---

[54] See my §2.1. Cf. Van Cleve (1999, 20–21) for a different view—that Kant's logical criterion is essentially the same as Frege's, Quine's, and even Carnap's.

However, one overarching point of this chapter has been to show that understanding the role which different conceptions of analyticity play is, partly, understanding the projects and ambitions of those who employ it; thus, misattributing technical conceptions to a thinker is not as innocent as it appears.

At any rate, the important feature of "Frege-analyticity"—i.e., the claim that a sentence is analytic iff reducible to the form of a logical truth by the substitution of synonym for synonym— is that it specifies a semantical criterion. That is, the definition relies on the interpretations of what words mean. This is significant for Quine's attack because he argues that this notion of analyticity relies on a problematically circular group of intensional, and thus semantic, concepts. An "extension" is said to be the class of objects that fall under a given concept; e.g., the general concept "football" has as its extension the set of all footballs. An "intension", by contrast, is meant to be the rule we might follow to pick out extensions—the rule in virtue of which we discriminate footballs *as* footballs.[55] An intensional notion, then, is one which concerns not extensions and referents directly, but how they are picked out: "meaning", "concept", "proposition", "synonymy", and so forth. As Quine notes, synonymy between two expressions cannot be cashed out in terms of which objects fall under them, for the same objects fall under "creature with a kidney" as "creature with a heart" and yet the terms are not synonymous; thus, he concludes, synonymy is not extensional (1951, 21).

The next move in Quine's argument, then, is to demand that we clarify synonymy (since the definition of analyticity invokes it) and to do so in a way which is extensional. Quine thought that the natural attempt to clarify the notion of synonymy would be to invoke the notion of definition. For, by *definition* "unmarried men" is substitutable for "bachelors". But, he then objected, depending on how we understand "definition", we get ourselves into further trouble. For, on the one hand, "definition" may mean the entry in the dictionary we look up, but if so this would actually presuppose synonymy, since the lexicographer plausibly first finds pre-existent facts about synonymies and then compiles them. Or, on the other hand, a "definition" may be something we explicitly stipulate or legislate ("This term: '. . .', will henceforth mean this object: . . ."). However, although for Quine this stipulative case counts as legitimate analyticity, he believes this to be limited to few scenarios—"would that all species of synonymy were as intelligible" (1951, 26).

---

[55] Incidentally, that rule might not be so easy to specify (which, to my mind, also affects what the extension is); for all intents and purposes, a discarded aluminium can in many parts of the world, within a game, is most definitely a football, as would be the conventional football on the shelves of a sports shop or in a factory. Then there is the complication that such a ball might not *count* as a football when thrown onto the pitch during a game in which players already possess and play with a ball; the latter is in a sense eligible—it is *the* football.

Quine hopes, then, to demonstrate that subsequent attempts to clarify the notion of synonymy fail too—for instance, if we try to use "interchangeable", with caveats, Quine reasons we must slip into an intensional "necessarily" (1951, 29–31). Accordingly, taking these points all together, Quine levies a circularity charge at those who wield an analytic-synthetic distinction, insofar as analyticity goes beyond "logically true". That is, according to Quine, it is impossible to define "analyticity" without reference to intensional notions, which each stand in need (he thinks) of extensional clarification. Thus, by Quine's lights, the notion of analyticity suffers two fatal problems: it assumes an untenable (because atomic) picture of the relationship between language and the world, and it suffers a vagueness which it cannot clarify except by use of notions themselves unclear.

There are many points of criticism against these arguments which we could rehearse.[56] However, as mentioned, the contemporary mistrust of analyticity does not have so much its source in Quine's arguments as in the atmosphere which those arguments brought in their wake. Perhaps that is not entirely fair; it is the meaning-scepticism which most subsequent attacks sought not to reproduce. For example, Williamson's own criticism, as we discuss in Chapter 3, invokes an epistemological holism, which naturally has a Quinean shape. And both Boghossian and Williamson, respectively an advocate and critic of analyticity, subscribe to a view which I think is at least latent in Quine as well: that descriptions and prescriptions are not logically exclusive. Indeed, come Chapter 4 we see that Quine's own conception of how language or theory and the world relate is implicit in the way he thinks analyticity must be, and this has terrible consequences for our justificatory place in the empirical world.

My concern in the following chapter is to scrutinise Boghossian's present-day attempt to recover analyticity from a weakened state by carving off its apparently more problematic aspects into a separate, unviable conception, leaving behind a workable notion. I argue that this approach is instructively wrong-headed.

---

[56] For example, see Grice and Strawson (1956) for a direct response, as well as Putnam (1962) for a critical discussion. More recently, see Bealer (1987) and BonJour (1997).

Chapter Two
The normativity of analytic truth

In the first chapter, we surveyed several conceptions of analyticity or associated analytic-synthetic distinctions. There are three main results to take from the chapter and to keep in mind throughout the rest of this thesis.

First, the distinction is pivotal to various philosophical projects or systems, often in being used to help articulate the nature or domain of philosophy itself, or some special part of philosophy. For example, Kant thought the future of metaphysics lay with the question how synthetic a priori judgments are possible. Or consider Carnap, who at first construed epistemology as construction of the scientific language, which includes showing which of the language's sentences are analytic or synthetic. Later, although he thought epistemology itself should be replaced by what he called logical syntax, or the logic of science, the logicist constructional project remained, which again required distinguishing the logical from the empirical parts of some framework language by appeal to analyticity (or L-truth, its explicated concept). There is no question, then, that historically the analytic-synthetic distinction has been fundamental to major areas of philosophical thought.

Second, as can be gleaned again from the difference between Kant and Carnap, conceptions of analyticity are not equivalent. Even though Frege (1884, §3) saw himself as utilising Kant's notion, and Carnap (1947, §2) as (later) explicating those of Kant and Leibniz, rarely do criteria for analyticity coincide neatly across projects. For example, despite much surface similarity, Leibniz and Kant discuss conceptual containment in quite different terms: for Leibniz, not only are some predicate notions contained within subject notions (which for Kant are only the analytic ones), *all* predicate notions are contained in their relevant subject notions, meaning that for Leibniz and not Kant all true affirmative propositions are analytic (though not all necessary). Or consider how the criterial role of logical non-contradiction in Kant and Leibniz differs to Hume (and is not explicitly part of Frege's or Carnap's conceptions): in Hume, the emphasis is predominantly psychological, depicting the impossibility of our being able to imagine contrary ideas at once, such as "red" and "green"; Kant, by contrast, employs a logical principle, and uses it to cash out the idea of conceptual containment. Furthermore, Kant applies an epistemological criterion absent from the other thinkers' conceptions as discussed, since a judgement is synthetic, for Kant, whenever it increases our knowledge (i.e., is ampliative, not explicative), although Anderson argues this is a consequence of the containment criterion and not a criterion itself (2015, §1.2.3).

Mention of Hume's psychological principle brings to mind the third and final point: if we fail to take account of the wider philosophical projects within which conceptions of analyticity feature,

and by which they are motivated, then we have an attenuated and often misguided interpretation of them. For there is good reason, as I argued, not to take Hume's epistemological fork as a version of the analytic-synthetic distinction at all, once we understand that the Humean "ideas" between which he thinks certain a priori relations hold lack the discursive nature of Kantian "concepts". That is, ideas are not, in the manner of Kant's concepts, applied by the understanding to empirical sense-data, thereby rendering perceptions intelligible; rather, they *are* such sense-data, intelligible as is, and mentally processed via perception (in Kant's terms, by sensibility). Accordingly, Hume's fork ought to be understood as part of his empiricism, and when taken as such, we have good reason to doubt its reputation as an incipient analytic-synthetic distinction. Moreover, even distinctions explicitly cast between the analytic and synthetic call on us to examine the theoretical detail which surrounds them. For example, the common catchphrase "truth in virtue of meaning" is not an especially helpful gloss on Carnap's differing conceptions of analyticity without at least a glimpse into his evolving logicism. And so, to make a fair appraisal of some conception, it is imperative that we examine its surrounding philosophical project and purpose.

Now, having surveyed these different historical conceptions of analyticity, we might compare them in terms of their most prominent criteria or aspects. Thus, we could say that Leibniz had an ontological conception, Kant an epistemological conception, and Carnap a logical conception. I do not think this would be an especially fruitful or indeed accurate means of description, except as perhaps a useful shorthand. Yet contemporary discussion is framed in terms of just such a general division—specifically, a divide is proposed between two apparently distinct "metaphysical" and "epistemic" conceptions of analyticity. Immediately this seems odd.

In §§1.3-1.4, after exploring Boghossian's discussion, I argue that in fact the distinction between so-called metaphysical and epistemic analyticity is, historically speaking, unwarranted. More problematically, however, it is badly conceived. I show that the attempt to distinguish metaphysical from epistemic analyticity fails to take seriously a fundamental point about the logic of our language: that rules are not, and cannot be, descriptions. To argue this point, I begin in §2 by examining a famous example from Wittgenstein's *Philosophical Investigations* (1953)—a discussion on whether we should or should not say that a standard metre is itself a metre in length. I use this example because Paul Boghossian (1996, 1997)[57] himself cites a Kripkean interpretation of it to motivate the claim that some rules also describe facts.[58]

---

[57] Boghossian (1996) contains the core arguments, but since Boghossian (1997) contained the same material, with expansions, I will henceforth reference this later article.

[58] For what it is worth, even though Boghossian is in some respects loyal to Kripke's work, Kripke himself describes analyticity in a way that does not cohere with Boghossian's story of separated metaphysical and

However, I find this attempt unsuccessful though instructive, and turn in §§3-4 to examining more generally why norms may not be descriptions. This involves discussing analytic truth with a normative focus that thus far I have not emphasised, and responding to an objection which opens up a pragmatic way of thinking about normativity. Finally, in §5, I bring the narrative back to the supposed division between metaphysical and epistemic analyticity; with arguments at hand for the normativity of analytic truth, we can see exactly why contemporary attempts to carve it up along these lines will never do justice to the idea of specifically conceptual truth.

§1 Boghossian

§1.1 Two kinds of analyticity?

We have seen in the first chapter that, similar to Kant, the way Frege defines analytic truths is partly epistemological. Though no reference is made by Frege in his definition of analyticity to how we do in fact know an analytic truth, the notion of an "ultimate ground" for an analytic truth is epistemic. (We should not forget though that the criterion Frege applies is as much logical as epistemological; e.g., "[i]f, in carrying out this process [of proof], we come only on general logical laws and on definitions, then the truth is an analytic one" (Frege 1884, §3).) Let us look in more detail at why we might care that Frege's definition of analyticity is wholly or partly epistemological in nature.

The recent tradition which splits the characteristics of analyticity into metaphysical and epistemological conceptions goes back to Boghossian (1997). His concern is to defend a notion of analyticity against Quine's arguments, because he thinks analyticity allows us to explain the possibility of a priori knowledge—particularly, for his purposes, the apriority of sentences which implicitly define the logical constants. Boghossian agrees with the main thrust of Quine's arguments against analyticity (Boghossian 1997, 335); however, he thinks that they apply not to analyticity *simpliciter*, but a particular, "metaphysical" conception of it. Before we discuss these terms, I note that Boghossian is right to suggest that there are different conceptions of analyticity, but he splits them only into two general camps—as the first chapter showed us, we might add "conceptual", "logical", and "semantic" to the list too, *if* we are in the game of naming a conception after its most prominent aspect.

---

epistemic conceptions. Though brief and sketchy, an analytic truth for Kripke (1972, 122f, n. 63) is both a priori and necessary, and so has (respectively) both epistemic and metaphysical aspects.

But that is just the point. As I soon argue, these labels are only emphases: traditionally, thinkers who stress the epistemological aspect of their conception of analyticity are still committed to metaphysical, or non-metaphysical, views. Indeed, for reasons I make clearer later, Boghossian's attempt to divorce aspects of any given conception of analyticity into distinct conceptions in themselves fails. Moreover, seeing analyticity aright shows us that there is indeed something interesting, on a metaphysical level, about a plausible notion of analyticity, though, as hinted, that detail is more aptly described as *non*-metaphysical. Before I can offer these arguments, however, we must first get to grips with Boghossian's terminology and aims.

Boghossian introduces these two conceptions—metaphysical and epistemological—by outlining two motivations for the introduction of an analytic-synthetic distinction, especially as conceived by the Vienna Circle. The first motivation he cites is that of explaining how a priori knowledge about the world is possible—e.g., how we can know logical and mathematical statements, as well as conceptual truths. The second of the Circle's motivations, Boghossian claims, is the problem of explaining *de re,* a priori necessities: i.e., how there can be necessary, language-independent connections, known independently of experience. I do not want to dwell much on these attributed motivations, but we can briefly reflect on the first, given the discussion in the last chapter.

As we have seen, the empiricist leanings of the Viennese group certainly made them concerned about such a priori systems as mathematics, logic, and metaphysics, which they sought to explain (or else obviate). Indeed, they were not satisfied with their predecessors' attempts to account for the a priori. We have already seen that Kant appealed partly to a priori intuitions to explain our knowledge of reality. Further, he argued that while our access to mathematics is a priori, mathematical judgements are synthetic. Accordingly, Moritz Schlick made it a key point to deny the possibility of synthetic a priori judgements (1925, 74; also §38). And although Frege, as per his logicism, saw arithmetic as analytic because reducible to logic, we discussed evidence to suggest that Frege believed in logical intuition (see Kanterian 2012, 25ff). Moreover, he still appealed to spatial intuition to explain geometry, which he saw as synthetic a priori, but post Hilbert, naturally, the Viennese saw geometry as analytic (Schlick 1925, §38). So certainly the Circle had something approaching the first motivation Boghossian attributes to them, namely that of explaining a priori knowledge.

Anyway, it is against these two worries—the apriority of mathematics and logic, and the metaphysical status of necessities—that Boghossian wants us to see how the different conceptions arose. First, then, the epistemological conception of analyticity is introduced, on Boghossian's account, to explain the apriority worry: "a statement is [epistemically analytic] provided that *grasp*

of its meaning alone suffices for *justified belief* in its truth" (1997, 334—emphases added). Take a candidate analytic sentence, such as "All adults are grown-ups", and consider how it is known. A sufficiently competent English speaker need not consult surveys of adults to find out whether they are predominantly, or all, grown-ups. No: one only has to know the meaning of "adult" and "grown-up" (and, of course, the logical vocabulary) in order to be justified in believing that the sentence is true (or so many claim[59])—this is what makes the analyticity involved here epistemological, according to Boghossian.

Next, Boghossian claims the metaphysical conception of analyticity is introduced to explain away the strangeness of *de re* necessities. On this conception, a statement is true by virtue of its meaning "provided that, in some appropriate sense, it *owes its truth-value completely to its meaning*, and not at all to 'the facts.' " (1997, 334). Thus, on the so-called metaphysical conception, no (important) reference is made to how the truth is known, but rather how it is "made true"—this, I presume, is what makes the conception metaphysical. (Though, as we shall shortly see, I regard this as a misdescription.)

As a side note, I think we should already be concerned about Boghossian's story here. As discussed previously, it is not clear at all that there were distinct motivations (here, metaphysical; there, epistemological) which, as it were, fully guided thinkers either to metaphysical or epistemological conceptions. It is likewise unclear that individual philosophers employing some notion of analyticity were not simultaneously committed both to its metaphysical and epistemological characteristics. Boghossian's account of the history, then, raises a preliminary doubt about the coherence of the distinction to which he ultimately cleaves. (Glüer (2003) also casts doubt, *en passant*, about the historical accuracy of Boghossian's logical empiricist survey.)

Nonetheless, having articulated the different conceptions of analyticity, Boghossian then argues in favour of epistemological analyticity and against metaphysical analyticity. Accordingly, Boghossian asserts that we can explain the possibility of a priori knowledge—specifically our knowledge of certain basic sentences which define the logical constants—without committing ourselves to explaining away *de re* necessities. Only in §1.3 will I argue that Boghossian's distinction between kinds of analyticity does not bear critical scrutiny. At present, I rehearse his positive argument for how we can explain a priori knowledge by use of this so-called epistemic conception.

---

[59] Timothy Williamson (2007, chap.4) dedicates a whole chapter to arguing against epistemological conceptions of analyticity, partly in response to Boghossian, having spent the prior chapter arguing against metaphysical conceptions. My next chapter critically overviews and responds to both these sets of arguments.

To begin his explanation of epistemic analyticity, Boghossian focuses on what he terms "Frege-analyticity": a statement is analytic if and only if it is reducible to a logical truth by the substitution of synonyms for synonyms. This particular notion of analyticity is one which Quine attacks in "Two Dogmas" (1951), but it in fact goes back past Kant to Leibniz (Glock 2003, 154; Burge 2003b, 200); the caveat with the Leibniz attribution is that, in his writings, this process of analysis is not something finite minds can undertake in the case of contingent truths, which require a so-called infinite analysis.

Its basic application is this: an analytic sentence is one which is reducible to the form of a logical truth, e.g.,

$$(\forall x)\ (Ax \wedge Gx) \rightarrow Gx$$

by the substitution of synonyms for synonyms—thus, "All adults are grown-ups" reduces to "All grown-ups are grown-ups" (given that "grown-up" is synonymous with "adult").

Boghossian takes Frege-analyticity to specify a "semantical condition", because it explains how analytic sentences are known by appeal to something about their meaning; this, he thinks, could be the only general method of explaining epistemic analyticity (1997, 337). Accordingly, Frege-analyticity aims to explain, on Boghossian's reading, how a cognising subject could come to know the truth of certain statements by recognising and substituting synonyms in the appropriate way.[60] However, Boghossian worries that appealing to Frege-analyticity in order to explain how a priori claims are epistemically analytic encounters an obstacle. Namely, that logical truth—specifically the elementary laws of first-order logic—would also need to be a priori for the explanation to work (1997, 345–46).

According to Boghossian, Frege simply assumed the apriority of logic in response to this worry. In fact, this is incorrect, given (as I pointed out in the first chapter) that Frege held a truth to be a priori if it is derivable exclusively from general laws. Thus, for Frege, the basic axioms of his system (general logical laws such as the law of identity) could neither be analytic nor a priori; rather, they are self-evident, which is sufficient in Frege's view for their epistemological purpose. Nevertheless, I suspect Boghossian would be just as unsatisfied with this response as he is with the response he attributes to Frege (cf. Boghossian 1997, 337–38). For derivability from self-evident general logical

---

[60] Boghossian (2003, 18f) offers a step-by-step guide for how someone might run through such a justification for their knowledge of a Frege-analytic statement, which he calls the "synonymy template".

laws would be too narrow a definition of apriority for Boghossian, whose aim is to explain how we can grasp logical truths independently of experience.

Indeed, Boghossian argues that Frege-analyticity ultimately provides an incomplete explanation, because it fails to explain precisely the analyticity of logical truths themselves. Moreover, he points out, Frege-analyticity cannot explain the analyticity of "Whatever is red all over is not green" and other conceptual truths, since they are not amenable to this synonym-substituting procedure (1997, 338–39). Such sentences he regards as "Carnap-analytic".[61] And it is a Carnapian or Wittgensteinian solution which, Boghossian claims, can provide the requisite semantical condition to explain how any sentence is epistemically analytic. Although Boghossian's aim is ostensibly to explain epistemic analyticity *tout court*, as already hinted, Boghossian thinks any semantical condition we put forward will reduce down to logical truth, such that we must explain the a priori grasp of logic if we are to explain how any proposition is epistemically analytic (1997, 345).

By Boghossian's lights, the semantical condition which can explain the apriority of logic, and eventually epistemic analyticity, is "Implicit Definition". This thesis[62] maintains that logical constants acquire their particular meaning via our stipulating that certain sentences involving them are true.

> *Implicit definition:* It is by arbitrarily stipulating that certain sentences of logic are to be true, or that certain inferences are to be valid, that we attach a meaning to the logical constants. More specifically, a particular constant means that logical object, if any, which would make valid a specified set of sentences and/or inferences involving it. (Boghossian 1997, 348)

Kathrin Glüer (2003, 44) points out that this process of implicit definition is unlikely to be the origin story for any given piece of logical vocabulary; on her recommendation, then, we should interpret Boghossian as spelling out how "and", etc. *could* come to mean what they do. (This does, however, prompt a question over why we should be concerned over some alternative way that words acquire meaning, and not their actual way.) In any case, Boghossian's statement of the thesis is somewhat obscure, and I shortly explore some of its issues; a clearer expression is offered by Bob Hale and Crispin Wright:

---

[61] According to Hacker, it is precisely such propositions that the Circle struggled with because they were not analytic, and not reducible to logical truths (1996, 49).

[62] That is, "Implicit Definition" as a thesis which Boghossian lays out. Of course, the notion has its roots at least in Hilbert's work on geometry, which we do not have space to consider.

We take some sentence containing—in the simplest case—just one hitherto unexplained expression. We stipulate that this sentence is to count as true. The effect is somehow to bring it about that the unexplained expression acquires a meaning of such a kind that a true thought is indeed expressed by the sentence—a thought which we understand and moreover know to be true, without incurring any further epistemological responsibility, just in virtue of the stipulation. (Hale and Wright 2000, 288)

Let us take an example. On this view of implicit definition, the negation symbol,

¬

comes to mean "*negation*" because a certain sentence (or set of sentences), such as

¬ (p ^ ¬ p)

is said to be true.

Inevitably, then, there is a pre-existing notion in the picture, though the symbol being defined does not yet call it to mind; the pre-existing notion (in Boghossian's language, "object") just has particular logical properties (e.g., that it negates). Accordingly, stipulating the truth of some canonical sentence involving the symbol—supposing this is an intelligible process, prior to interpreting one of the symbols—casts light on which notion the symbol must depict, given the logical properties in play. The thesis claims that by the stipulative move, I attain at once an a priori knowledge of the sentence's truth and knowledge of its meaning; moreover, I know the sentence is true given only this meaning-knowledge.

A natural way of reading this stipulating scenario is that by making the stipulation I am reminded of the meaning of some other symbol—that, as it were, the meaning which I already know is made transparent to me and I just see that it must carry over to the new symbol. However, Hale and Wright argue otherwise—that I need to think of myself as "so far innocent of the conceptual resource which the implicit definition affords" (2000, 300). Accordingly, the implicit definition account of logical constants is meant to show that anyone really could grasp the meaning of logical constants through this procedure.

The reason Boghossian invokes this implicit definition thesis is to explain the apriority of logical truth, our knowledge of which in turn is supposed to explain how epistemic analyticity is possible. Let us see how he proposes implicit definition suffices to explain a priori grasp of logical truth in general; we can call it the Boghossian argument (hereafter, BOG):

1. If logical constant C is to mean what it does, then argument-form A has to be valid, for C means whatever logical object in fact makes A valid.

2. *C* means what it does.

Therefore,

3. *A* is valid.

(Boghossian 1997, 348)

Before making this less abstract, observe BOG's logical structure. I suppose Boghossian intends the argument to have the following form:

1. P → R
2. P
Therefore,
3. R

Yet really there is a third proposition, conjoined with the first to provide a conjoined antecedent, so that to (P) "logical constant *C* is to mean what it does" is added (Q) "*C* in fact means whatever logical object makes *A* valid". (I sweep aside, here, issues we ought to raise over just what is meant by an object, logical or otherwise, "making" an argument-form valid.)

Accordingly, we now have the following form:

1. (P ^ Q) → R
2. P
Therefore,
3. R

Yet, if that is BOG's form, it is invalid. For (Q) needs to be included in the second premise; i.e., for Boghossian to conclude that some argument-form *A* is valid, he needs it to be the case both that *C* means what it does *and* that *C* makes *A* valid. (This problem recurs in the following exposition.)

With the more formal concerns out the way, let us try to make BOG (more) concrete. To interpret its first premise, imagine an argument-form in which the constant *C* appears, which implicitly defines that constant; *C* is stipulated to mean whatever makes that argument-form valid. Suppose we try inserting the conjunction symbol into the following form,

p
Therefore,
p *C* q

such that we get

p
Therefore,
p ^ q

The latter form is clearly invalid, given the meaning of conjunction (a meaning which we do not yet know, when making the implicit definition). So when *C* "makes the form of argument valid", according to Boghossian's thesis of implicit definition, it is because the meaning of *C* is disjunction and not conjunction.

The second premise of BOG—that *C* means what it does—is meant to be fairly straightforward. (Though, in an appendix, Boghossian (1997, 358–62) deals at length with the worry that this premise is unknowable a priori.) I confess, however, to being unsure as to what it *really* asserts, for lack of seeing a sensible contrary here. Presumably, it is impossible (or worse, simply meaningless) for a term not to mean what it means, at least insofar as we think of "term" in a conceptual and not merely orthographical manner. I am unsure, therefore, that BOG's second premise is any more informative than "*C* has a meaning". Indeed, to recall the above point about logical form, Boghossian here needs (Q)—i.e., "*C* means what makes *A* valid". Without that, and if (2) is ultimately uninformative, in turn we ought to doubt (1), since its antecedent contains the phrasing "if *C* is to mean what it does"—as opposed, exactly, to what? I will not follow up this issue here, and simply jot it down as notational awkwardness on Boghossian's part. And for the time being, at least, let us presume BOG is valid.

Boghossian's application of the implicit definition thesis in BOG is supposed to explain how we have an a priori grasp of sentences which implicitly define the logical constants, which are logical truths, by consideration only of their meaning; in turn, this is to explain the idea of epistemic analyticity. For example, take some candidate analytic sentence, the kind which Boghossian calls a conceptual truth: "All adults are grown-ups". Boghossian thinks this is Frege-analytic, where that means it is reducible to the form of a logical truth by substitution of synonym for synonym; e.g., again

$$(\forall x)\, (Ax \wedge Gx) \to Gx$$

And if BOG is valid, it can apparently then be deployed to show that this logical truth is knowable a priori, since we can know the meaning of the logical vocabulary involved simply by stipulating as true the sentences implicitly defining *them*. Thus, only by means of stipulation and our awareness of meanings are we entitled to knowledge of an analytic truth, which all remains an a priori exercise. Incidentally, it remains utterly obscure how this argument helps us to explain the analyticity of other conceptual truths to which Boghossian gives lip-service, such as colour-exclusion propositions—"Whatever is red all over is not green" (Thomas Kroedel (2012, 147) raises the

same concern). This is something else we shall have to brush over, in order to cover Boghossian's more fundamental points, at which point we encounter a significant problem.

The first main stage of Boghossian's argument is thus over. The second is his denial that the implicit definition thesis has two apparently troubling consequences: (either or both of) "conventionalism" and "logical non-factualism". In §1.4, I contend that Boghossian's argument to this effect is badly misconceived. But in order to grasp that point, we shall first have to expose another (related) mistake, exhibited by Boghossian's refusal of the non-metaphysical characteristic of analyticity. In this respect, what we soon see is that Boghossian's argument recalls an old critique of analyticity, and this basic argument itself fails to motivate the divorce of (non-)metaphysical from the epistemological aspects of any notion of analyticity. Instead, as the remaining sections aim to show, the argument to which Boghossian appeals makes a fundamental error in assuming that rules may describe or report facts.

§1.3 Analyticity and the non-metaphysical

The first thing to note about the so-called metaphysical conception of analyticity is that Boghossian's terminology is misleading (though it has, to be sure, caught on[63]). It is not clear how the conception is *metaphysical*; it seems, rather, that the logical empiricists to whom Boghossian credits the notion had an *anti*-metaphysical aim: to show that analytic statements were not made true by some worldly state of affairs. Viewed this way, the metaphysical conception of analyticity is non-metaphysical. To avoid that obscurity or potential confusion, I will use Burge's suggested terminology henceforth—the "vacuousness" conception of analyticity (2003b, 200).[64] This label helps articulate the idea that analytic statements are true "come what may", in a non-factual manner; an adult is a grown-up, and no changes in the world could have affected this.

Boghossian thinks that arguments derived from Quine successfully show the vacuousness conception to be untenable. He attempts to undermine the conception ultimately by appeal to a principle he thinks we all ought to accept; as it turns out, this principle has a fairly long history. It only seems to be hinted at by Quine when, in "Carnap and Logical Truth", he motivates the notion of (logical) truth in virtue of meaning by noting "that sheer verbal usage is in general a major determinant of truth" (1962, 351). That is, Quine thinks, every sentence is true partly in virtue of

---

[63] See, e.g., Williamson (2007), Russell (2007), and Papineau (2015).

[64] Künne opts to call the vacuousness conception the "autarky" conception, which denotes self-sufficiency: when an analytic sentence is said to be true in virtue of its meaning, it thereby (in some way) secures its own truth; it is self-sufficient (2006, 216).

what the words in it mean (for, the thought goes, if they meant something different, a different claim would be made). And so the parallel thought, which is only implicit in Quine's 1962 article,[65] is that the world is *also*, in general, a major determinant of truth. (This "two-factor" picture of language and world is something we pick apart, with the help of Davidson and McDowell, in Chapter 4.)

Eventually Roderick Chisholm (1966, 37–38) took Quine as his cue when he argued that though truths about the empirical world are true partly in virtue of meaning, analytic sentences are *also* true partly in virtue of the facts.[66] For example, Chisholm suggests that the English sentence "Being round excludes being square" is true "if and only if, being square excludes being round." Thus, since the right-side of the bi-conditional refers "to a relationship among properties and not to rules of language or ways in which we use words," the sentence "cannot be said to be true solely in virtue of the ways in which we use words" (1966, 38). Boghossian's point is very similar to Chisholm's:[67]

> 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"? How could the mere fact that S means that p make it the case that S is true? Doesn't it also have to be the case that p? (Boghossian 1997, 335)

So convinced is Boghossian by this point that he refers to it subsequently as a "meaning-truth truism" (1997, 336). The putative truism holds that no sentence is true solely in virtue of what words mean, or how we use them—rather, every sentence must be true because of what the words mean and something factual, or "how things are" (Williamson 2007, 59). I contend that Boghossian's reliance on this much-invoked principle begs the question against anyone who does not view the notion of a priori necessities that exist "in the world" or otherwise factually as intelligible to begin with.

For, one should press, what fact about the world helps make true a sentence such as "Adults are grown-ups"? *Pace* Mill, we know that this sentence is not used on the basis of an inductive generalisation. So we ought to wonder how the facts are arranged to "make" this sentence true. Boghossian has in mind retorts such as this: adults are grown-ups by definition, and the truth of

---

[65] Initially written in 1954.

[66] Glock (2003, 159) and Schroeder (2009, 100) source the point back to C. I. Lewis.

[67] To my knowledge, Boghossian (1997) is not aware of Chisholm, though his incredulity at the vacuousness conception is expressed in strikingly similar terms: "What could this possibly mean?" (Chisholm 1966, 37); "What could it possibly mean to say that…" (Boghossian 1997, 335).

"Adults are grown-ups" depends on the self-identical nature of everything in the world (Boghossian 1997, 335). The problem with this retort is that it presumes what it sets out to prove, *again*. It presupposes that there is some substantial, worldly property called "self-identity", and it asserts that every item in the world has this property as a matter of fact. But this is just to say, of what may equally be regarded as a rule—e.g., that one cannot sensibly deny that something is itself—that "Adults are grown-ups" reports a fact. And the factual status of rules is the very point at issue! (It may turn out that rules are descriptive after all, but one could not hope to demonstrate this by restating the point.)

In any case, Boghossian wields the above principle against the idea that any declarative sentence is capable of being vacuously true or non-factual. So we should at least stop to ask how this vacuousness conception comes about, according to Boghossian's story. He has it that the logical empiricists were right to focus on the explanatory role that could be played by the idea of implicit definition—specifically, on its explaining some a priori knowledge. However, he believes they fatally took implicit definition to have two consequences: conventionalism and (logical) non-factualism. (Essentially, then, the vacuousness conception is defined by its embrace of these two principles.)

In Boghossian's terms, "conventionalism" holds that analytic sentences are made true by convention, or definition, while "non-factualism" maintains that such sentences do not express claims. Accordingly, Boghossian's rejection of the vacuousness conception consists in attempting to show that implicit definition entails neither of these theses. Thus, he thinks, one can make an epistemological point about how we know certain basic logical truths without thereby committing to some non-metaphysical account of logical truth. This is how, Boghossian argues, the analytic theory of the a priori can be "salvaged from the wreckage of the linguistic theory of necessity" (1997, 332). (I admit to finding Boghossian's partitioning of these two "consequences" itself rather baffling. For, as I will argue over the course of this chapter, when a proposition is conventional (i.e., normative), it is necessarily not descriptive—prescribing and describing are distinct sorts of logical move.)

§1.4 Conventionalism and non-factualism

Let us hear more about so-called conventionalism and non-factualism, then. Boghossian is not concerned to identify any particular representative of these positions; he is opposed generally to any thesis according to which some statements are both true by convention and not made true by the facts. I am sure, nonetheless, it will not harm us to take a brief look at how the conventional and non-factual character of analytic truth features in Carnap's work, and how Wittgenstein

characterised logical truth, since these are the figures to whom Boghossian attributes the implicit definition thesis: "I want to worry about the fact that neither Carnap nor Wittgenstein was content . . . with Implicit Definition. Typically, both writers went on to embrace some form of irrealism about logic" (1997, 348).

Propounded in his later phase, Carnap's definition of an analytic sentence as one that holds in every state-description—as true, corresponding to every row of a truth table—is a nice gloss on the idea that analytic sentences are true come what may, no matter how things stand in the world of facts. And we saw in the prior chapter that the logical empiricists used convention to explain mathematics. Partly influenced by Hilbert, and partly by Wittgenstein's *Tractatus* (1921), Carnap and others held that analytic sentences were tautologies, because consequences of truth-tabular definitions, which are conventions of symbolism. For example, in the *Aufbau* (1928), Carnap regards a theorem as analytic when it is deducible from the conventional axioms of a constructional system; that is, an analytic truth follows from a definition, which is conventional. So, indeed, Carnap's conceptions of analyticity were generally characterised as non-factual and conventional.

The idea that logical truths follow from arbitrary conventions of symbolism is what the Circle apparently took from Wittgenstein's *Tractatus*. In this work, Wittgenstein had argued that logical propositions were tautologous, if not contradictory, as could be displayed in truth-tabular notation. Yet, as Baker and Hacker discuss (2009, 356–70), Wittgenstein did not see logical propositions as seriously arbitrary, for such propositions, he held (at this stage of his thinking), "flow from the *essence* of the elementary proposition – not from arbitrary conventions, but from the essential bipolarity of the elementary proposition" (2009, 362). That is, Wittgenstein regarded the most basic proposition, which "asserts the existence of a state of affairs" (1921, §4.21), as essentially true or false, and this is because it depicts an atomic fact, which must either obtain or fail to obtain, depending on the arrangement of the world. Given this essential "bipolarity" of an atomic proposition, the idea of negation, he held, lies latent within it—for the elementary proposition must be either true or false; a proposition's being false is equivalent to its being negated. In a like manner, Wittgenstein argues that all of logic emerges from the logical structure of this most elementary proposition. And of this structure, only its signs are to be regarded as conventional:

> logic is not a field in which *we* express what we wish with the help of signs, but rather one in which the nature of the absolutely necessary signs speaks for itself. (Wittgenstein 1921, §6.124)

So much for the idea that Wittgenstein saw logical truth as conventional, in the manner of the logical empiricists. But what about its non-factual character?

Wittgenstein indeed saw logical propositions as non-factual in the *Tractatus*, and this is something that remained in his later philosophy. Although, in §6.124, he says that logical propositions "describe the scaffolding of the world", he clarifies this as a representing. In the context of the *Tractatus*, the idea must be that they "show" the necessary structure of the world, while not being descriptions: they are strictly speaking "about" nothing; their constituents lack sense (cf. Hacker 1986, 53). And it is precisely this view which separated him from, *inter alia*, Russell, who held that logical constants were the names of logical objects (Baker and Hacker 2009, 311). Thus, when Boghossian writes into his idea of implicit definition that a constant, *C*, "means whatever *logical object* in fact makes [its argument-form] *A* valid" (1997, 348—emphases added), given that he thinks Wittgenstein would accept "logical objects" into his ontology, it is not surprising that he is puzzled about why Wittgenstein would think logical propositions are non-factual. At first, Wittgenstein is committed to the idea that logical propositions, being tautologous and strictly saying nothing, have no sense. By his middle period, Wittgenstein argues that they do indeed have sense, which is constituted by their formal rules. It is this idea which entices Boghossian, and indeed prompts him to cite Wittgenstein on the matter:[68]

> It looks as if one could *infer* from the meaning of negation that "~~p" means "p". As if the rules for the negation sign *follow from* the nature of negation. So that in a certain sense there is first of all negation, and then the rules of grammar. (Wittgenstein 1974, 53)

(Note the "as if".) And

> [w]e would like to say: "Negation has the property that when it is doubled it yields an affirmation," But the rule doesn't give a further description of negation, it constitutes negation. (Wittgenstein 1974, 52)

Boghossian wants to retain this constitutive view of the negation operator, but to divorce it from Carnap's conventionalism and early Wittgenstein's non-factualism. (Later Wittgenstein does not think logical propositions are factual either, but not for the same reason: he accepts that they are either normative or systematically related to inference rules; see Baker and Hacker (2009, 252), and my §3.) Yet even in this very citation, Wittgenstein urges the reader away from seeing the rule of

---

[68] Boghossian (1997, 348) cites these passages as though they follow one another, but this is incorrect—the latter passage precedes the former by a number of paragraphs.

negation as *describing* anything, precisely because it constitutes negation; a rule for a symbol is not descriptive because it is constitutive.

As we have seen, Boghossian hopes to show that his attempt to explain epistemic analyticity—i.e., that some sentences are knowable in virtue of their meaning alone—via arguing that we can grasp meaning of logical constants via implicit definition does not entail either that analytic sentences are conventional or non-factual. How does he hope to do so? Largely (and fatally), he relies on a Kripkean take on later Wittgenstein's "standard metre" example, which we examine in §2. Some preliminary marks are in order first.

Boghossian accepts that logical truth is at least partly conventional:

> according to Implicit Definition, "If, then", for example, comes to mean the conditional precisely by my assigning the truth-value True to certain basic sentences involving it; for example, to "If, if p then q, and p, then q". And in an important sense, my assigning this sentence the value True is arbitrary. (Boghossian 1997, 351)

As he continues, the decision to say that the "implicit definer" is true (I prefer to say that modus ponens reasoning is valid) is what gives the expression "if, then" the meaning that it does, such that prior to that decision, it simply did not have that meaning (or its meaning was incomplete). Yet, he argues, from our (apparently) choosing to make this sentence true it does not follow that conventionalism is true. Against the conventionalist, Boghossian argues that conventions only determine what a sentence expresses and not the facts which make what is expressed true.

At the outset this looks confused. Consider Boghossian's example sentence, "Snow is white".[69] He thinks that by "assigning the truth value 'true'" to the sentence, we determine what it expresses (*that snow is white*), but that we do not thereby determine snow's being white. Something else has to do that, he reasonably opines: what makes snow white has nothing to do with our linguistic conventions. Quite. But look again: determining snow's being white is just another way of determining that "Snow is white" is true. But apparently we can "assign" truth to the sentence; so which is it? There is a contradiction here.

The whole line of thought, in my view, is tangled. The argument may initially seem compelling because it draws on the now influential principle we discussed above, as offered by Chisholm (and others). Here is that principle simplified:

> (P) For any given true sentence S, S is true because it expresses some proposition P, and P is true.

---

[69] He also cautions this sentence may not be an implicit definer, but presumably thinks it inessential to the point (Boghossian 1997, 352). See also Rundle (1979, 451) and Russell (2007, 57).

Applying (P) to some substitution instances renders some examples. "Snow is white" is said to express *Snow is white*, and, of course, snow *is* white. Likewise, "Adults are grown-ups" is said to express *Adults are grown-ups*, and adults *are* grown-ups.

However, the blanket principle (P), despite its wider acceptance, is deeply problematic. For it conceals the actual function of supposed analytic sentences under the innocent-looking term "sentence", which is necessarily mute to the function of sentences in practice. I argue that we must attend to this special function if we are to understand what it means when we say that such sentences are true (see esp. §5). This function is *normative*, which is necessarily not descriptive.

That point has been recognised, historically, to differing degrees. We saw in Ch. 1 that Frege acknowledged the normativity of the most general logical laws, though it seems for him they still described a Platonic realm of logical objects, with which we are intuitionally acquainted. However, in the *Begriffsschrift* Frege occupies a curious position in regards specifically to definitions of new signs (1879, §55). He acknowledges them as prescriptions which are therefore not yet judgements, but once the definition has been made, they are "transformed" into analytic judgements (which, it seems, makes them descriptive). Normativity is in the frame at this point, then, but it did not extend to analytic judgements. More relevant still is Carnap, whom we saw, at least by his pluralist period, did not think of analytic sentences as factually true, since we are free to choose different logics. Likewise, the early Wittgenstein certainly thought that the propositions of logic had no content (though they do show or represent the scaffolding of the world; see my Ch. 1, §2.2 on these points). Indeed, Alberto Coffa, whose account of Wittgenstein and Carnap inspires Boghossian's implicit definition thesis, claims exactly this: that since sentences which implicitly define their constituents function as *rules* which we are "free to endorse", they thereby do not "tell us anything that is the case" (Coffa 1991, 266; cited in Boghossian 1997, 350).

Yet this point—that logical propositions are rules without factual content—is exactly what Boghossian rejects:

> I don't see that there is any inconsistency between supposing that a given logical principle – for instance, the law of excluded middle – serves to implicitly define an ingredient logical constant, and supposing that that very sentence expresses a factual statement capable of genuine truth and falsity. (Boghossian 1997, 350)

It turns out that the whole weight of Boghossian's argument against the vacuousness (i.e., non-factual) conception of analyticity is placed on this rejection. Accordingly, it must be probed further. I side here with Coffa and against Boghossian: a sentence cannot simultaneously be a rule (e.g., a definition) and a description of the facts. In what follows, we examine the sentence on which Boghossian focuses in his attempt to show to the contrary that a rule can describe a fact.

§2 The standard metre

§2.1 Abstract standards

Boghossian argues that his implicit definition thesis entails neither conventionalism nor what he terms non-factualism—the thesis that analytic sentences do not report facts. To make his case, Boghossian cites Kripke's reflections on the "standard metre in Paris" scenario discussed in Wittgenstein's *Philosophical Investigations* (1953, §50). There, Wittgenstein argues that one object can neither be described as one metre long, nor described as not one metre long, and that is the standard in Paris which was then used to determine what is to count as a metre.

Kripke, in response, thinks Wittgenstein "must be wrong" (Kripke 1972, 54). After offering his reasoning (which we scrutinise shortly), Kripke goes on to reveal that the sentence "S [the standard metre in Paris] is one meter long at time t_0 [the time of its stipulation]" is both contingent and a priori (1972, 54). It is contingent, he claims, because the stick used to define "one metre" could have actually not been a metre: it is a contingent matter of fact that the stick S was the length that it actually was at time t_0—it could, for example, have had heat applied to it, and stretched. And he thinks it is a priori because it can be known at the point of stipulation (for example, by the person who stipulates it) without any further investigation (1972, 55).[70] Thus, in one fell swoop Kripke appears to have severed the fork between apriority, analyticity, and necessity on the one hand, and aposteriority, syntheticity, and contingency on the other. So there is much at stake: if investigation may be a priori but still issue in contingent truths about the world, then armchair reasoning is much more powerful than it prima facie seems (cf. Schroeder 2006, 245).[71]

In order to make his case that implicit definition entails neither conventionalism nor non-factualism, Boghossian needs it to be possible that individual sentences may at once describe facts about the world *and* function as (expressions of) norms. With Kripke as his support, this is indeed what Boghossian asserts: "there appears to be no inconsistency whatsoever between claiming that a given sentence serves to implicitly define an ingredient term and claiming that that very sentence expresses something factual" (1997, 350). So it is crucial that Kripke's reading of the standard metre case is intelligible, and further, that it justifiably leads to Boghossian's conclusion.

---

[70] See Evans (1979) for a rebuttal. Evans claims that Kripke only shows the claim to be "superficially" contingent—a feature of how it embeds inside modal operators—which does not entail it is "deeply" contingent, such that its truth would depend on some state of affairs.

[71] Albeit slightly less powerful than on Kant's view that philosophy can deliver synthetic a priori truths. For since necessity is a criterion of apriority, for Kant, critical philosophy can tell us truths which not only go beyond "mere" concepts, but which tell us about the necessary structure of experience.

But Kripke missed the point of Wittgenstein's remark. Following his claim that the standard metre in Paris is describable neither as one metre nor as not one metre in length, Wittgenstein remarks that "this is, of course, not to ascribe any remarkable property to it, but only to mark its peculiar role in the game of measuring with a metre-rule" (1953, §50). Kripke responds confusedly:

> This seems to be a very "extraordinary property", actually, for any stick to have. I think he must be wrong. If the stick is a stick, for example, 39.37 inches long (I assume we have some different standard for inches), why isn't it one meter long? Anyway, let's suppose that he is wrong and that the stick is one meter long. (Kripke 1972, 54)

This is far too quick. As Baker and Hacker have argued, focussing on the role played by the stick in practice allows us to see where Kripke goes wrong (2005, 195ff). The standard metre bar in Paris, when it was indeed used as the standard, acted as a *canonical* sample. That is, though other metre bars measured against the bar in Paris (or just outside Paris, initially, in Sèvres) are standards against which to measure the length of various objects and spaces, the bar in Paris itself is the one to which all such metric standards are subject.[72] The moment at which the bar is given this canonical status is the moment at which the phrase "one metre" should call forth "the length of the standard metre bar in Paris".[73] Prior to this, of course, the phrase "one metre" would not have that significance. To be historically accurate, it should have called forth "the length of the 'mètre des Archives'," but *this* metre as a provisional standard was subject to problematic variations which the international standard in Sèvres was designed to overcome. (We might say, insofar as our purposes go, the original meaning was vague or ill-defined—sometimes it is good for our terms to be vague; for standards of measurement, it is very often a bad thing.)

I have just claimed that the phrase "one metre" only receives its particular meaning at the point of stipulation. It is consistent with this claim that "one metre" had a different meaning prior to the stipulation, just not the same one as at the point of stipulation. Or we can simplify the issue by thinking of a fictional scenario. If the British standards authorities were so inclined, they could ratify a new unit, the "Metryard", by use of a canonical sample. Then the phrase qua unit of measurement would only have a meaning at the point of stipulation, since *ex hypothesi* it had none before. Therefore, it is easy enough to see that stipulated terms and phrases receive their meaning *at*, and not before, the moment of their stipulation.

---

[72] The distinction between canonical and standard samples is also Baker and Hacker's (2005, 192), though, as we will see, it does map fairly accurately onto the historical development and dissemination of the standard metre.

[73] See §2.2 below for a sketch on how this normativity, in practice, loses its bite and spreads out diffusely.

Now, the reason it is significant that the phrase "one metre" only means what it does at the point of stipulation is that while the bar soon to become the standard metre obviously has a length, it does not, prior to attaining its status as the measure for other items, have the length of *one metre*. It might seem that Kripke (and Boghossian with him) would grant as much. But actually, given that Kripke thinks the phrase "one metre" has a referent (namely, *the length*), which the bar itself happens to instantiate—which we use as an attempt to "fix" the reference of the phrase "one metre"—presumably, he cannot even grant this much (1972, 55). So, for Kripke, even before the bar in Paris is stipulated to be the canonical sample for "one metre" (before it is defined as such), it is, if a good example (and how would we know?), already a metre long. But this is not fully intelligible. For it so radically breaks with how we understand and operate with measurements. One is left to wonder how we would ever know whether a standard matches up to its abstract form. We would need, *per impossibile,* to hold the stick alongside the abstract, non-physical form.

Let us return, again, to see how the unit of measurement is used and designated, how it functions in practice; for there we come to realise that an abstract length is of precisely no significance. Baker and Hacker broach these points by way of noting that the standard metre bar became not an example (of the length), but the *sample* by means of which the length is to be determined (2005, 195). If we ask what the bar is supposed to be an example of, and we are told "*the length of one metre*", it becomes plain that (insofar as this is intelligible), what the stick is an example of is completely out of our (material) hands. But this simply fails to accord with the facts: as recounted, the bar decided in Sèvres to function as the canonical standard for measuring one metre replaced the mètre des Archives formerly so used, and it was decided that the new bar should match the exact length of the mètre des Archives in whatever state it then currently was in. That is, owing to the unreliable nature of the former measure, the new standard was not designed to pick out an ideal which the last one tried and failed to pick out; rather, the new standard, for ease and relevance, was determined to be the same length as the last one because, with better construction and standardisation, the new stick would do a better job at being canonical than the old one.

§2.2 Breaking standards

There is one important caveat to this discussion, which deserves attention. There is perhaps a legitimate sense in which we may describe the standard metre bar as not in fact being one metre long, and indeed such that this description has normative weight. Any standard which loses its place must do so because it is not deemed fit for purpose. The purpose for which a metre standard would be deemed unfit is measuring; if its length changes according to pressure or temperature, it

is not up to the task. Imagine the conversation had by the appropriate board of measurement authorities. It is plausible they would remark, "The stick is not a metre today—look! Something must be done about this." Now, is this a completely illicit remark? One would be tempted to say so given the apparently strict Wittgensteinian view that one must not predicate "of the thing what lies in the mode of representation" (1953, §104). But in that case we run into a tension, because such remarks are surely intelligible, and play a motivating role in the changes of standards.

My response to this problem is to urge, first, that we see the nature of Wittgenstein's example as an idealisation of a point. That point remains firm: if there is one and only one sample of x, against which another item y, is measured, then while y is in fact degrees of x-ish, x itself cannot be. My second response is to note how the standard metre case, in practice, likely enters into greater complexity. For consider that the length designated by the standard metre has, as it were, trickled down from such places as Sèvres into general population familiarity. There is a tangible sense as to what counts as a metre in day-to-day dealings with lengths. Not least because our world is fairly rigid, so far as we can tell: our bodies and other objects are not constantly in flux, changing lengths (Baker and Hacker 2005, 191); in part the everyday objects of the world become standards of each other (Wittgenstein 1956, §§93–94); it is all very familiar. Occasionally, standards are appealed to in order to settle disputes about lengths. Suppose then (as is true) that many metre sticks in circulation are regularly being used for accuracy and to settle arguments; what then if the canonical standard starts to change its length (cf. Wittgenstein 1956, §5)? The measuring practices would continue, and within an intact normative structure. It is likely we would hold up against the canonical sample our standards and note that the sample is not really a metre in length. I think what this would demonstrate is not the looseness of our talk, but how normativity in many cases is itself much more diffuse and pliant than we might think. The authority of the performance at stipulation does not stretch forever into the distance; rather, it is in part divested and spread as communities begin to authorise their own standards. But this makes for a much less punchy example.

In any case, with this detour, we can affirm that prior to stipulation the bar shortly to become the canonical standard metre bar was not one metre in length. And we can maintain this, without contradiction, even if we note that it matched the length of the then-current standard for a metre; for when it matched up to the metre standard, the bar itself indeed was a metre, but that was prior to its becoming the standard. All the same, cannot Kripke still maintain that *after* stipulation the bar is one metre in length? One route to maintaining this has now been blocked; i.e., the attempt to maintain that the metre after stipulation is a metre in length for the reason that the metre already was a metre in length. So, in order to maintain this anew, another route has to be found. And

presumably, that route can only begin with the stipulation itself. That is, something must be significant about *the moment* where it is decided, ratified, that this particular bar is to serve as the canonical sample of one metre, with which to measure any derivative standards of a metre. Indeed, paying attention to what this moment achieves sheds light on the whole problem.

§2.3 Stipulating standards

As we saw in §1.4, Boghossian suggested that "if, then" only gets its meaning when certain basic sentences in which it figures (such as the tautologous expression of modus ponens)—the implicit definers—are called "true". But he then professes his puzzlement at the fact that both Wittgenstein and Carnap apparently thought the sentences used to implicitly define a term are not factually true, or are *only* true by convention (Boghossian 1997, 348). In order to maintain that a thesis of implicit definition (for example, about the logical constants) does not entail such a "logical non-factualism" or conventionalism, Boghossian appeals to the just rehearsed argument from Kripke against Wittgenstein regarding the standard metre.

Here, I want to point out that it is *prima facie* inconsistent for Boghossian to cite Kripke in defence of this claim, for Kripke, as we just discussed, maintains that the standard metre *before* stipulation already either is or is not a metre in length. By contrast, presumably Boghossian has to take the same tack as with "if, then", and is committed to saying that prior to stipulation the term "one metre" did not have a meaning. Instead, he needs to say that the stipulation itself is what gives the phrase its meaning. This reveals the difficulty that Boghossian gives himself by offering a Kripkean defence. For we now have a situation where Boghossian and I agree that, contra Kripke, the stipulation made in France that *that* platinum-iridium bar is to count as the standard metre is what now gives the phrase "one metre" its meaning, yet we disagree over what this means about the sentence which expresses that stipulative definition, "This stick is one metre long". It is worth, then, exploring the argument in a little more detail.

Boghossian asserts that a sentence which implicitly defines an ingredient term may, consistently, also express a fact about the world (1997, 350). Why does Boghossian think this would be consistent? Presumably he has something like the following scenario in mind. The bar itself prior to stipulation as the standard metre of course has a length which can be measured by some existent unit (if there is one), but it cannot yet be one metre unless we are, like Kripke, Platonists about length (and then we have other problems to face). So we must ask what is special about the (act of) stipulation itself which allows that particular bar to be described as a metre in length. I take it Boghossian, and many others, would think this: there is nothing "special" about that moment, other than the stipulation (and presumably other requirements) setting into play the system of

metric measurement; without that moment, we have no warrant to call anything "a metre". But after that moment, as a trivial matter, any measurable item whatsoever, including the standard metre, is susceptible of being either one metre or not in length.

I agree entirely with the claim that any measurable item whatsoever is susceptible of being either one metre or not in length after "one metre" is defined as "the length of the standard metre bar". It just depends on how we understand "measurable". Are tables measurable by metre rules? They certainly are. Are metre rules measurable by metre rules? Indeed, sometimes. Are stretches of gloopy, thick oil measurable by metre rules? Possibly not—it depends. What, then, is the sense in which the standard metre bar is not measurable? It is clearly not the sense in which thick gloopy oil cannot be measured by a stick. To resolve this, we need to notice that it is not just the *item being measured* which determines whether it is measurable, it is also the *item doing the measuring*. And not just in terms of the physical suitability of the measuring tool, but sometimes for a deeper, non-physical reason. That is, the standard metre bar, though perfectly able to measure other items, is necessarily unable to measure itself: this is a logical feature of its role in our measuring practices.

In the *Investigations*, Wittgenstein wonders about a strange scenario: "Imagine someone saying, 'But I know how tall I am!' and laying his hand on top of his head to indicate it!" (1953, §279). The remark is humorous because such a person, if he intends to be taken seriously, has failed to tell us anything. Certainly, if the situation had called on him to demonstrate his knowledge of heights, or his height in particular, he has resolutely failed. He may as well say that he is as tall as himself. His case is analogous to the standard metre case. If we want to maintain that the definition of the phrase "one metre" entails that the stick used in that ostensive definition is itself one metre, we are committed to saying that this stick is "as long as itself". Well, insofar as that is an intelligible form of words here, it is certainly not informative.[74] A builder measuring a doorway would not accept from his assistant the answer that the doorway is as wide as it is without perhaps presuming his assistant was indicating it was time for a break.

We can throw light on this situation by finding the main source of ambiguity: "is". Schroeder argues that Kripke performs a "conceptual sleight of hand" which exploits precisely the ambiguity that Wittgenstein sought to show his reader with the example (Schroeder 2006, 246). "This stick is one metre long" can mean either "This stick is (as a matter of fact) one metre long" or "This stick is (to be called) one metre long". As Schroeder explains, when Wittgenstein says that the metre bar in Paris is not a metre long, he means to deny the former, factual sense of the copula, not (what Schroeder calls) the "analytic", latter sense. Hence the sense to the claim that the

---

[74] Cf. Evans (1979, 182–83).

standard metre in Paris is not (in fact) one metre long, but is (to be called) one metre long (cf. Wittgenstein 1974, 53–54). And if one wants to maintain that the factual and definitional senses of the copula are unified, one only has to consider the man who is as tall as himself: what *fact* is conveyed by saying that "this stick is as long as this stick", where both occurrences of "this" point to the same exact stick? (Schroeder makes a similar point in a later piece (2009, 96) in connection with the truth "A young swan is a young swan"—it gives us no information.)

In response to this question, one could not cite, say, a general logical truth about identity as the fact conveyed, without begging the question (this is the point I made in §1.3 against Boghossian's Russellian retort). That is because one would, as a prior matter, have to establish that logical truths were facts and not norms, or that what expresses a norm may also simultaneously describe a fact.

§3 Why norms may not be descriptions

In the previous section, I showed that a particular sentence token—"Stick S is a meter long at time t_0"—may never be, on one occasion of its utterance, both descriptive and normative. This is helpful for my broader aim because it undermines Boghossian's positive argument for a coherently distinct epistemic conception of analyticity. Still, Boghossian may retort, I have not ruled out *any* sentence or proposition as simultaneously normative and descriptive; rather, I have only focused on propositions about canonical standards. I tackle that challenge in this section by progressively expanding my case. First, I show that the class of sentences which can only express either (but not both) descriptive or normative propositions is wider than that containing only canonical standard-measure sentences. This offers preliminary support for the thought that no sentence may express a proposition with a dual descriptive and normative character. Following this point, I show that, in fact, it is impossible for any sentence to play both these normative and descriptive roles at once.

Work on the normativity of propositions today largely finds its roots in one of two sources: the later philosophy of Wittgenstein or Wilfrid Sellars. For example, Glock (2003), Baker and Hacker (2005, 2009), and Schroeder (2006, 2009) fall into the first camp; the "Pittsburgh" school[75] comprising (alongside Sellars) Brandom and McDowell, and, more generally, *global* inferentialists[76] fall into the second camp—e.g., Brandom (1994, 2001), Lance and O'Leary-Hawthorne (1997), and Whiting (2007). There is a fair amount of explicit and implicit crossover between these two

---

[75] For a useful introductory overview of the tenets broadly shared by these thinkers, see Maher (2012).

[76] Murzi and Steinberger (2016) distinguish between those inferentialists who focus on some local area (typically logical), such as Boghossian, and those who apply inferentialism globally to the whole of language, such as Brandom.

camps; for example, very commonly, the notion of an "inference ticket" is employed in writings of both groups. (The notion of an inference ticket—a license that allows one to infer one proposition on the basis of another—was coined by Gilbert Ryle (1949, 105).) I focus on the Wittgensteinian strain here for ease of exposition.

Baker and Hacker (2009) discuss Wittgenstein's work on what he called "grammatical" propositions—those traditionally classed by philosophers as necessary. The range of necessary propositions is made up of different types: mathematical, logical, and "metaphysical". Examples such as "2+2=4", tautologies, and "Nothing is red and green all over" fall respectively into those types. Baker and Hacker maintain that necessary propositions are normative and not (unless used to make different moves) descriptive; moreover, they hold that their normative status makes them "concept-forming" and "partly constitutive of the meanings of their constituent terms" (2009, 259). They also point out that the distinction between grammatical (necessary) sentences and empirical ones is not to be understood as a distinction between sentence types, but sentence tokens. That is, what determines that a given sentence expresses a grammatical proposition instead of an empirical one will not be the form that the sentence takes—declarative, let us say—but rather, its particular uses. The authors claim, e.g., that the sentence "Water is (= consists of) H_2O" was once used as an empirical conjecture but later hardened into a rule (2009, 249), and so now plays a role in concept-determination; i.e., the concept "water" is now partly determined by the rule (2009, 260).[77] The first of these claims—that a particular use of a sentence may not be both descriptive and normative—is prima facie not the most intuitive. At least insofar as Boghossian is concerned:

> 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)

This is not intuitive to me (quite the opposite), but clearly the claim needs further support. Not least, also, because it is the fulcrum on which my larger argument is balanced. In explaining the idea themselves, Baker and Hacker use the example of a rule of chess, and in what follows I play with the example further.

---

[77] Baker and Hacker offer this claim in order partly to disabuse the reader of the notion of the necessary a posteriori, popularised in contemporary writings by Kripke and Putnam. For its status as a rule, as will be seen, is claimed to block its status as a description of a fact, such that on the *one* occasion of utterance it cannot be the sort of thing we only come to know through experience (a posteriori—with caveats) yet which simultaneously attains a necessary status (because it describes the way the world *must* be). See Baker and Hacker (2009, 259–60).

Consider "The chess king moves one square at a time" (2009, 263). This sentence could be used to describe, empirically, the happenings on a chess board. Imagine a chess examiner (I do not suppose such characters exist) whose role is to overview a set of games and ensure that the chess kings are used only to move one square at a time (except when castling). We can imagine that she utters the sentence as a description of the occurrences she sees, and ticks her box accordingly. Consider next, however, the sentence as used in a rather different way, as written in the chess manual, or dictated by a chess-teacher. The sentence "The chess king moves one square at a time" written or spoken in those circumstances does not describe empirical happenings, and does not correspond to a fact; there is no chess board, king, or player in mind, or in the relevant context.

Indeed, for pedagogic purposes it would be catastrophic if there were particular facts to which rules corresponded—for teaching the roles of chess pieces is to be generic, thereby allowing a person to play the game at an unspecified range of boards. It is distinctive of most manuals that they are not only historical; it is distinctive of many helpful beginners' manuals that they are, further, ahistorical. A questionable manual it would be, if a manual at all, which mostly described practices rather than prescribing how effectively to practise. This is not to say that one cannot read off the rules of a practice from that practice; to the contrary, this *must* be possible, and is the main point argued for in Ryle's (1949, chap.2) work on knowing-how and knowing-that. My point is only that manuals principally help those who are struggling to pick up rules simply from what they experience; they can expedite this normative learning.

Here we might object that it is, nevertheless, an empirical matter of fact that people *do* use the king to move one square at a time (or so the chess examiner hopes), and so the rule is simply a description of the fact that within chess communities, the king is to move only one square at a time. This is a view similar to that held by A. J. Ayer, in *Language, Truth, and Logic* (1936), where he saw analytic propositions as a subset of empirical propositions which describe how we use symbols.[78] He later abandoned the view, recognising it as a mistake to equate empirical with analytic (and in his view, a priori) propositions. As Glock notes, "[o]ne must distinguish between the expression of a rule and an empirical proposition to the effect that a community follows certain linguistic rules" (2003, 162). That is, statements about which rules are active do not place obligations on anyone, as those rules do: if I recount the rule "In England, the red traffic light signals cars, lorries, and bicycles to stop", I have not thereby told my audience to stop in such circumstances—though given antecedent knowledge of rule-following and laws, they should infer that they ought to stop when at a red light in England.

---

[78] This is pointed out in Glock's (2003) discussion.

Keeping in mind, however, that normativity applies to sentence tokens, one must remember to attend to the function of the sentence in discourse. Were I to utter, as a frustrated driving instructor, "In England, the red traffic light signals cars, lorries, and bicycles to stop" to my speed-hungry trainee at the wheel, I may well be placing a normative constraint on him by sardonically expressing the rule. Taken out of context, my words may seem to form a description of a section of the highway code. However, anyone who responded in such circumstances by continuing to run through red lights, simply acknowledging "yes, that is the meaning of the red light in England", is either being rude or has taken grammatical surface structure too seriously and thereby failed to take the normative hint.

This much shows, I think, that sentences normally expressing rules of chess and of the highway code may have descriptive uses in addition to their standard normative uses, and that grammatical form is not a foolproof guide to the real purport of a sentence. Further, we have seen that the sentences in focus here, when functioning as rules, do not describe any circumstances, so they share this feature with standard-measure sentences. All the same, we have not conclusively shown that a sentence on one occasion of utterance absolutely *cannot* be normative and descriptive simultaneously. This negative, general claim is one for which Baker and Hacker argue; to do so, they urge us always to question what possible state of affairs or facts a given rule is supposed to describe. As we saw in the above example, a chess rule no longer expresses a rule if it is used to refer to particular chess boards and pieces (unlike, e.g., an explanation of such a rule). Rules are prohibitions, licenses, guides; these seem to require generic application, and it is not clear how a description—a sentence actually used to describe (and not simply a declarative sentence)—can ever play such roles.

For a contrast, consider an apparently metaphysical proposition: "Nothing can be red all over and green all over". Baker and Hacker observe that this proposition is often thought to express an internal relation between the natures or essences of red and green (2009, 264ff); or, if it is different, that it describes the necessary structure of the world. They draw on Wittgenstein, however, to show it rather acts as an "inference ticket" (2009, 263), licensing us to infer, e.g., that if this chair is green all over, therefore it is not also red. We do, with our descriptions, present the world as being a certain way, and we do after offering such a presentation often "travel" from that presentation to another. However, we license ourselves to do so not by the use of further descriptions on which we fall back, but rather with rules or norms for "*re*-presenting" the way the world is (2009, 250). I am justified or licensed to describe a man as a bachelor, having presented him first as unmarried, because of the norm of representation, "A bachelor is an unmarried man" (2009, 263).

It seems to me there is a crossover here with those thinkers, such as Sellars, who argue that there is a form of inference which appears enthymematic, and so formally invalid, but is yet valid in a "material" sense (Sellars 1953, 313). E.g., consider the following argument:

> The ball is red.
> Therefore, the ball is coloured.

A natural instinct (natural for many philosophers, in any event) is to say that the argument is missing an implicit premise; informally phrased, something such as "Whatever is red is coloured". Sellars argues, however, that the argument as displayed is already valid, prior to inserting an apparently missing premise. His claim is that the missing premise is in fact a rule of inference—a rule that concerns the content or "matter" of the concepts in play, which is thus a "material rule of inference". The rule in our example is of course a famously discussed truth of the kind that puzzled Carnap, and which Boghossian calls "Carnap-analytic". It is the paradigmatic example of a grammatical proposition as discussed by Baker and Hacker, and, as we just saw, it is the sort of proposition they see as a norm of representation, allowing us to travel from one premise to another. For Sellars, the rule is implicit because it is a propriety already embedded in linguistic practice (as norms of representation are supposed to be also); it is a part of the content of the employed concepts, and therefore plays a crucial rule in constituting conceptual content (1953, 334).[79]

Likewise, Baker and Hacker are at pains to emphasise that looking into the normative role of necessary truths reveals their concept-determining role. The chess rule "The king moves one square at a time" is a rule at all because, indeed, it is constitutive of a chess king that it moves one square at a time (2009, 263). "Red is darker than pink" is again partly determinative of the concepts "red" and "pink", and therefore the proposition entitles me to say of an object which is red that it is darker than another which is pink (2009, 263–64). However, one still wants to object further: is it not the case that red is, *in fact*, darker than pink, such that I can legitimately use the sentence as a description of the world—indeed, of the necessary structure of the world? Baker and Hacker's own tactic for explaining disanalogies between obviously empirical propositions (say, "London is cold in December") and necessary propositions is to show how we treat them differently when we ascribe truth or (more significantly) falsity to them. We are invited to notice that understanding an

---

[79] For further discussion see Brandom (1994, 100ff), though we shall see in Chapter 4 that Brandom's inferentialist understanding of Sellars provides an ultimately untenable account of linguistic meaning.

empirical proposition (a description) partly consists[80] in knowing "what is the case if it is true and what is the case if it is false" (2009, 273). Yet, of course, this cannot carry over to "Red is darker than pink"—for how can we conceive of a world in which red is not darker than pink? That is the point of its being necessary.

As a side note, I am sure there are contexts in which it is apt to describe bright reds as lighter than dark pinks, so this example likely needs adjustment—either in the form of a clause about paradigm cases, places on a colour spectrum, or a total change. In the lattermost case, I have in mind "Red is closer to orange than it is to yellow". Hume uses such an example as one of his relations of ideas, which as discussed in Chapter 1 are psychologically conceived: i.e., according to Hume, we cannot imagine the ideas obtaining in a manner other than the relation depicts them. In what follows, I think we can use this Humean example in place of Baker and Hacker's favoured proposition, without much trouble. In Chapter 5, however, we shall have occasion to doubt the rigidity even of such propositions; regrettably, this is a complication we cannot enter into at present. Suffice it to say that given my argument here, grammatical propositions do not reflect the necessary structure of the world, and so a change in what is properly grammatical does not entail a change in how things must be; it does not reflect some kind of metaphysical instability or inconsistency.

Indeed, rules change. Recall an earlier example which Baker and Hacker suggest had "hardened" from being an empirical proposition into a rule: "Water is H_2O". We *can* imagine a world in which water is not H_2O; Putnam's (1975) famous thought-experiment trades on this imagined possibility. In my view, we ought to recognise that some necessary propositions, in virtue of having a traceable history back to their pre-normative employments, tempt us more easily into treating them as descriptive even while normative. See how this sentence is notably different to "No square is round". It seems to me there is a rough, though useful, distinction to be made here between a subset of necessary propositions which have had pre-normative employments, and a subset which have always (or very nearly always) had normative status.[81] "Water is H_2O" is employed now as a rule, and not as a description; its former non-normative status helps deceive us into thinking it also functions descriptively now. Having made this distinction, let us notice how

---

[80] Baker and Hacker do not say "consists entirely"; neither does he say "partly consists". I would wish to make the latter qualification to avoid the statement of a crude verificationism about understanding. Cf. Baker and Hacker (2009, 273f).

[81] The number of members in this latter subset might appear to be quite low, until one considers the huge range of propositions attaching, e.g., to new terms: "If I'm tweeting, then I'm sending a tweet on Twitter".

the apparent possibility of these dual descriptive-normative necessary truths tempts us to think, like Quine, that *all* necessary truths must partake of both descriptive and normative character:

> the lore of our fathers is a fabric of sentences [which] develops and changes, through more or less arbitrary and deliberate revisions and additions of our own. . . . It is a pale grey lore, black with fact and white with convention. But I have found no substantial reasons for concluding that there are any quite black threads in it, or any white ones. (Quine 1962, 374)

This is an attractive picture of the matter. But it is neither the only, nor the best one available. Indeed, as discussed above, the reasons for denying this picture are several and compelling.

Let us adjust our view: some type-sentences will, in their lifetime, begin as token descriptions but eventually express rules; some might go back the other way; but none of them will function at once both to describe and to prescribe. That is a fact of their logical role in our shared linguistic practices. Seeing things this way, we retain a certain insight from Quine: diachronically, the threads in our webs of belief are pale and grey. But *synchronically*, those threads are quite black and white.

§4 A pragmatic objection

I want to pause to take account of an underlying worry that practical matters call for a re-evaluation of this line of thought. The following example makes the case that the standard metre really is, in fact, a metre in length precisely because (and not, as I have been arguing, in spite) of the role it plays or could play in practice. So that, actually, it is possible for a proposition to be synchronically both part-rule and part-fact.

Imagine a scenario in which, during the *Great Measurement War*, the young "metric division" is under siege from the ageing "imperial division". During the imperial siege, the metric division find themselves cornered, with just their prime weapon left: a canon which only fires rods the length of a metre (and the width of a specified thickness—though other dimensions will not be in focus here, for simplicity). As their stock of sticks runs low, they strike a near fatal blow to the invading imperial forces. But it is of paramount importance that the metric division fire off their next stick urgently and accurately. Any nearby rods which are not exactly one metre in length will not suffice for the purpose. Their luck appears to have run out: there are no longer any metre sticks in their location.

Distraught, A has an idea:

> A: "Aha, we *do* have one metre stick left—the standard metre on the wall there, hermetically sealed!"

B: "But this will not work, because this stick is not a metre long; we are doomed."

A: "But surely it *is* a metre long. It's the same length as all the other sticks we've been firing. So it's a metre long. Now quickly—load the canon!"

B: "No, we really are doomed: it is not a metre long because it is the stick with which we *measure* or *determine* what is to count as a metre long in our measuring practices..."

A: "But—"

B: "...When you say, 'it is a metre long', you at best may mean, 'it is to be called a metre long', thereby reiterating a normative stipulation."

Anyone in these circumstances would think B's protestations rather silly. Indeed, in the moment, B's pedantry is dangerous. The consequences of something's being a metre long are many; in some cases, those consequences could be extremely important.[82] If my interpretation of the canonical standard metre proposition as having only normative force does not also bear those consequences, then the interpretation would run into serious trouble. Sticking resolutely to the line of thought I have been pushing would appear to obstruct sensible and practical aims. That sort of disconnect between thought and action is antithetical to the broadly pragmatic notion of analyticity I have presented thus far; a notion which emphasises the role of analytic sentences in practice.

So what are we to say about this case? Are we ultimately to concede that "This stick [the canonical standard] is in fact a metre long" is a true description? That would be hasty. What is at stake here is, first, my critique of Boghossian's distinction between epistemic and "metaphysical" (or, in my terminology, *vacuous*) analyticity—the thought that some sentences may be justified purely by virtue of their meaning while no sentences may be true solely because of meaning (or language, or convention). In turn, at stake is a vision of philosophy according to which the discipline articulates our socio-linguistic norms—and not, primarily, "facts about the world". Boghossian's attempt to carve off an epistemic conception of analyticity is, after all, an attempt to secure the role of a priori reflection about logic, without denying that those reflections issue in factual truths. Thus, for Boghossian, with respect to at least the subject matter of logic, philosophy takes a seat alongside other disciplines whose investigations lead to factual truths about subject matters. But if what I have maintained against Boghossian is correct—that analytic sentences are rules, and that rules may not be descriptions—then insofar as philosophy investigates analytic

---

[82] In the scenario we discuss, of course, the important consequence is rather contrived. But we do not have to look far for serious application—e.g., the importance of accurate and consistent measurements in the construction of skyscraper buildings.

sentences, or (if there is a difference) norms, it does not reveal, by a priori means, truths about a non-conceptual world. So much, then, is at stake.

Our response to this case must be twofold, in recognition of its idealised nature. First, we should treat the example as though it were not idealised; that is, to pretend, as in Wittgenstein's initial example (1953, §50), that there really is only one object to which we can compare metre sticks (namely, the hermetically sealed, canonical standard metre). The second response, however, shall be guided by the discussion in §2.2 and §3 above, in which we recognised the often diffuse nature of normativity in practice, and in turn the idealised nature of Wittgenstein's example. According to this more grounded understanding, the authority of a stipulation not only stretches but *fades* into the future, whereupon authority is taken up at different points in an intricate normative web. Let us now take the example at face value first.

§4.1 The idealised case

The problem of the canonical standard metre typically exhibits the following contradiction. Suppose a canonical standard, Stick S, determines what a metre is to be[83] (which, *ex hypothesi*, is not itself a metre in length). Then suppose that another rod, Stick S', matches up to S; thus, S' is a metre in length. Coupled with the premise that rods which match up to each other will have the same length, it seems we must conclude that S is also a metre in length, contra the first supposition. We can quickly resolve this contradiction by eliminating the first claim that S is not a metre in length. (This is what Kripke has done.) But there are other ways to go. Should we eliminate the premise that rods which match in length have the same length? That would be too great and unintuitive a burden to bear. Indeed, notice from the discussion in this chapter that neither I nor my interlocutors have denied that rods may have "the same length", nor the more basic thought that the canonical standard metre rod "has a length".

Incidentally, however, it seems as though Kripke thinks this is what Wittgenstein must have believed about the standard metre:

> Part of the problem which is bothering Wittgenstein is, of course, that this stick serves as a standard of length and so we can't attribute length to it. Be this as it may (well, it may not be) . . . (Kripke 1972, 54)

---

[83] It is this feature of the case—that one and only one item sets and maintains the standard—which makes it idealised, even though my treatment in this section asks the reader to imagine that this idealised case is the real case.

I have not found evidence that this is what Wittgenstein thought of the standard metre stick; it would be rather odd for him to think that the stick had no length at all. Indeed, for reasons given below, the stick's having a length is what made it apt for being a standard in the first place. (At least, in its historical context; for of course the contemporary standard defines a metre according to the distance light travels in a vacuum over a certain time.)

So we can argue, as I have, that the rod's status as a standard for a given unit prevents it from being measureable by that very same unit, without implying that it cannot be measured by another unit, or that more fundamentally it is not the sort of thing which has a length. All may consistently acknowledge, too, that anything which matches the length of S will "have the same length". Examining these two locutions—having a length and having the same length—will help us overcome the seeming contradiction here.

The contradiction generated by the standard metre case, relayed above, employs insights that we should expect anyone competent with "length discourse" to acknowledge: that both S and S' have lengths, and that any two rods $S^{n...i}$ which match in size have the same lengths. These platitudes about length do not ordinarily come into question; no-one is seriously perplexed about everyday, length-based transactions, which make use of or implicate such basic truths. Wire leads have lengths, brick walls have lengths; the idea of cheddar cheese does not have a length, and nor does anguish (though anguish may persist for a lengthy time). Since it is clear that all rods have a length (you might say it is analytic that they do), so it is clear too that S has a length—regardless of its being a standard of length for a particular unit. Accordingly, it is clear that no-one ought to resolve the contradiction presented by the standard metre case by means of argument that S has no length. Notice further that something which has a length, in virtue of its having a length, can be put to a number of purposes. *Inter alia*, it can fill gaps and holes that are themselves of "the same length" or shape, in the way that the idea of cheddar cheese or the phenomenon of anguish simply cannot (though, speaking less strictly, the idea of cheese may well "fit" the moment).

As with S's having a length, I likewise think one will not want to deny that S and S' are of the same lengths. If we examine talk about the sameness of lengths, we can actually find ourselves close to a solution. And it is not difficult to see what it is for two items to have the same length. Turning to just why we might require items to be of the same length is a useful cue here. The legs on a table should (if it is a normal table) be the same length so that they each stand perpendicular to a level floor and thereby support a surface parallel to the angle of that floor. Or the length of a timber, say, should be the same length as the distance between two walls so that the walls and timber provide mutual support to each other. In short, we tend to require items to be of the same length so that they can perform relevantly similar roles to each other—roles for which their length

is significant in the first instance. It is natural to surmise, following these remarks, that items have the same length insofar as they fulfil the same purpose for which having a length is significant. (Obviously this is rough, because items do not only have lengths but depths, widths, density, and more besides.)

Given this rough criterion for length, it would be incoherent to maintain that, as introduced in our example, even though S and S' both have lengths, they do not have the same length. In our example, the role for which it is relevant that an item has a length is to fill a particular canon. Of course S will fill that canon, because the canon gap is the same length as S; it was measured to be the same length! In general, it is intrinsic to canonical standards of lengths that they fit in the gaps they are used to measure—they would be rather useless otherwise (notwithstanding, as noted above, the non-physical standard used today). So neither pointing out that S and S' both have lengths, nor, further, that they have the same lengths, forces us to say that the length they both have is "one metre". (I said Kripke's argument was too quick; it is these points that he did not take into consideration.) All it urges us to do is to acknowledge what their having the same lengths amounts to; among other things, it amounts to the fact that both are suitable fodder for the metric canon.

One might still be unsatisfied: "The lengths are not the *same* if different things follow from them; e.g., if I can describe one rod truly as a metre in length, and the other not." The point of emphasis here is the culprit. Functioning as an adjective or pronoun, "same" attaches to a noun of a particular type—a person, idea, material object, etc. Even when attaching to one general category noun (a material object noun, say), further ambiguity persists: this apple is the same (colour, variety, taste, molecular structure) as that one; the book so-and-so mentioned is the same (copy, version, binding) as the book on the shelf. It is plausible that in ordinary scenarios, context implicitly sets the relevant aspect of sameness attributed to two or more items. But in philosophical cases, it normally helps resolve apparent paradox by making the hidden aspect(s) explicit. Consider: The first apple is the same as the second; the first when ingested makes one ill while the second does not; so they are not the same after all. Well, the first apple is a granny-smith—just as is the second—but the first had turned bad, so they were not qualitatively the same in respect of their nutritional aspects. Likewise, the lengths are the same stretches of matter, or distance between marking points, but they are not the same in units, for one length determines what the unit is to be, while the other does not.

Had A or B phrased their initial difficulty in terms of needing something to fill the metre canon, and not needing a stick a metre in length, then the solution would have been non-controversially apparent. But neither A nor B phrased their problem in those terms. Instead, A says that they do

have a stick a metre in length: the canonical standard. Given what we have uncovered above, we can see that really A speaks loosely—a looseness for which she is entirely forgiven (B's strictness, by contrast, in *these* circumstances manifests his own kind of, more worrying, incompetence). What A could have said is more cumbersome: "We don't have something a metre long but we do have something which determines the length of a metre and therefore fulfils the same purpose!" The latter formulation sacrifices usefulness for exactitude (and suffers for it). But ordinarily, this laborious phrasing would be unnecessary, owing to what should be a tacit awareness that the canonical standard is the same length as the other rods being used as fodder. That is, *of course* the standard metre rod will fit the canon, for it is the very stick after whose mould the fodder were designed! Likewise, B could have responded to A's claim by saying "It's not a metre, but it is the same length, so let's use it indeed". (However, maybe A would be best saving that comment until after the siege is over.)

In any case, the resolution of the apparent contradiction presented by this problem case has shown us that a standard of length must be the same length as the units made to its figure; indeed, that is a fact of everyday physics. But it has also kept true to the thought that we cannot describe the length of a standard and the lengths made in line with the standard, *jointly*, in the metric system.

§4.2 The grounded case

> Certainly the "definition" which is the lexicographer's report of an observed synonymy cannot be taken as the ground of the synonymy. (Quine 1951, 24)

As I pointed out in §3, the authority of stipulated norms does not work in practice in the clear-cut way suggested by Wittgenstein's idealised example. Stipulations can have, rather, a tendency to disperse and then weaken in their authority, travelling as they do across ever further plains of usage. Sometimes, too, as new variations on standards are formed, authority is sent back in the other direction: as Quine noted, lexicographers are faced primarily with the task of detecting those new standards. On Quine's view, judging the lexicographer as stipulating the meaning of terms would be to put the "cart before the horse" (1951, 24). However, this seems to me only half right: for after all, lexicographers solidify and codify the loose normative webs formed by usage, and thereby provide touchstones of authority for competent language-users and language-learners going forward. (The game of *Scrabble* is no fun without giving the full weight of authority to the dictionary.)

There is, then, an intricate back-and-forth between those who stipulate, articulate, follow, codify, and modify norms, and this whole process is normatively significant. In the passage

represented by the quote above, however, Quine is lured into searching (unsuccessfully) for the ground of synonymy, without which it is thought analyticity cannot gain a footing. If I am right, here and in many cases, it is folly to search for the ground to which norms can be reduced, unless the ground is something as very general as "practice". Contained within practice are those just noted contributors, who may interact—as when dealing with definitions and synonymy—in a non-hierarchical fashion; that is, neither exclusively top-down from legislators to rule-followers, nor solely bottom-up in the other direction. Accordingly, we should now wonder whether this model accurately depicts the standard metre case too.

In truth, there are few completely general claims to make about normative structures, such that we cannot assume what goes for one goes for all. For some rules are within systems relatively easy to change, with immediate and lasting effect—lawmakers sometimes have this ability, for example (though note how the system of precedents in common law legal jurisdictions allows for a bottom-up change of sorts). Centralised banks can change financial norms, as when they introduce or run-down a new unit of currency. Measurement norms, however, do indeed appear to be more like those governing words than those governing pennies. As we have seen, national standard metre rods were cut to the canonical international standard; from there, presumably regional and commercial sub-standards are cut to the national standards, and so on. As the standards spread out, each begins to reign over its respective jurisdiction, as it were. In this process, it is at some point inevitable that the notion of a metre becomes in a sense *abstract*.[84]

That is, participants in socio-linguistic measuring practices come to understand "one metre" as a rough-and-ready length of such-and-such distance. When they need to be more precise, they bring measured and measuring rods into play. But if a magician's rod were used for such a purpose, which subtly shrank when placed alongside objects, those in the practice would baulk at the authority of the stick. It would be less of a joke, of course, were all metre standards sold commercially to begin exhibiting the same shrinking features. For at that point, regular metre-stick-users—carpenters, let us say—are perfectly within their rights to proclaim: *that's* not a metre. How should carpenters go on, at this point? They could of course approach a shrinking canonical standard and remark *that's not a metre*, and still be warranted to make the claim. Given that the "force" of the canonical rod, in setting what counts as a metre, has dispersed, carpenters do not have to acquiesce to its authority when faced with it shrinking; they do not have to carry over a similar frailty to the length of one metre. It is not clear how they should go on. But the point is

---

[84] Note: *not* abstract in the sense on which Kripke relies when he claims that "one metre" is a rigid designator, having a referent—the length, *one metre*—which fixes the reference of the phrase; on this view, the standard metre rod itself is an example of that abstract length. See §2.1 above and Kripke (1972, 55).

this: it is a coherent dilemma they would face; they would have a certain amount of choice in the matter, and they could intelligibly describe the canonical standard metre as not in fact one metre long.

With these insights to hand, let us return to the *Great Measurement War*. If A and B work within a sociolinguistic practice much like ours, where the groundswell of usage plays a significant role in setting measurement norms, it is not necessarily wrong of A to claim that the canonical standard is in fact a metre long. On this less idealised interpretation, B's responses may not just be infelicitous: he is sticking too rigidly to an earlier articulation of a norm. The authority of the seniors in the metric division, over what counts as a metre, has waned for some time; the rod which is hermetically sealed is not the one which has final say—or at least in many circumstances would not be. The authority of the seniors is not that of a centralised bank when determining a new a unit of currency.

There are, thus, two ways to handle the pragmatic objection to the standard metre case. One way is to take the notion that there is one and only one criterion for what counts as a metre quite literally. When the challenge to employ the canonical standard metre in a practical situation arises, we can account for why it still functions in the way that other metre sticks do—why it *must* function that way. All rods have a length; all rods which match in length have the same length. These are more primitive truths than those articulated by conventions which stipulate what a particular unit-term is to designate by means of a certain sample of length. Those conventions do prohibit, however, our ability to talk about samples or criteria of units in terms of those very units themselves. Thus the standard metre, in the idealised case, is not itself a metre in length.

That point holds true of the more realistic way of handling the problem, too. On that interpretation, the then so-called canonical standard did not have the same authority it may once have had. Owing to how much traction the abstract notion of a metre has gained over time in public usage, it is no longer clear that the only recourse we have to what counts as a metre is the canonical standard. Perhaps the canonical standard could be reintroduced as the defining, unwavering standard among a sea of radically unreliable metres. In such a scenario, the first response to the pragmatic objection is once again directly relevant.

The relevance of both responses, however, cannot be underestimated. They allow us not to forsake basic truths about lengths, while retaining the thought that norms of measurement cannot, even in the most urgent circumstances, describe facts.

§5 Back to Boghossian and Kripke

It is time to return to our starting point. To recall, I had established previously that there is little reason to think that, historically, the conceptions of analyticity employed by philosophers in the past divide into two broad camps: epistemic and metaphysically vacuous. Philosophers such as Kant and Frege embedded their conceptions of analyticity within unique philosophical projects, and the character of those projects inevitably left its marks on how they conceived analyticity. I suggested that the best way, then, to frame these many different conceptions was in terms of emphasis or aspect: Frege emphasises the logical and epistemological character of analyticity, to suit his logicism; Kant focuses on the epistemological difference between synthetic and analytic judgements, in a bid to firm up the foundations of metaphysics, and so he also emphasised the explicative nature of analytic judgements, and the transcendental role of synthetic a priori judgements.

It is in the context of historical uses of "analyticity", especially by the logical empiricists, that Boghossian's distinction is supposed to have a home. Boghossian maintains that one conception of analyticity holds a sentence is analytic when it is true in virtue of meaning. Thus, "Adults are grown-ups" is analytic according to this conception, given the supposition that "adult" means the same as, or means part of, what "grown-up" means. The other conception is claimed to hold that a sentence is analytic when it is justified in virtue of meaning; i.e., when one can know it to be true only on account of knowing what its constituent terms mean. The motivating thought behind this epistemic conception is that regardless of how sentences turn out to be true, one can always know a subset of them (the analytic ones) on the basis of one's linguistic competence. Thus, also on this conception, "Adults are grown-ups" is analytic, because one's familiarity with what the sentence (and its component parts) means simply mandates one's knowledge of it. The claimed benefit to distinguishing these two aspects of analyticity is that one can ostensibly keep hold of the epistemological queerness of analyticity while rejecting the metaphysically vacuous consequences. That is, Boghossian wants to be able to say that basic sentences (those that implicitly define the logical constants) are known to be true by competent speakers on the basis of their knowledge of the logical language, which in turn explains their Frege-analytic knowledge of sentences such as "All bachelors are unmarried". And he wants to say this *without* maintaining also that strictly speaking such implicit definers are not factually true; that they do not correspond to some worldly fact.

The purported upshot of Boghossian's position, then, is that he can support the a priori methodology of philosophy, or logic, while evading certain criticisms historically and forcefully levied at analyticity. Namely, those criticisms which cast doubt on its vacuous, non-factual

character. For, Boghossian retorts, how could it be that even one sentence is true simply because of what it means? Surely, he thinks, every sentence is true on account, first, of its having some meaning, and, second, of that meaning's aligning with the way things are.[85] Boghossian is in good company in making this argument. As discussed, it has its roots in Quine, and later, Chisholm; moreover, critics of analyticity and conceptual truth more widely, such as Williamson, have adopted it too (see my Chapter 3, §3).

The problem is that this argument presupposes something false. It is not the case that all sentences are true in the canonical way imagined by Boghossian and others. (To anticipate an objection by Williamson, this is not the same as saying that there are different senses of "true"; we deal with this accusation in detail next chapter.) Of some sentences, to say that they are true is to reaffirm their status as norms, not to describe some situation that makes them true. Grammatical form can play a misleading role here, as with sentences such as "The chess king moves one square at a time", which is declarative in form. However, when one attends to the function of such sentences, when used on a particular occasion normatively, one finds that they do not purport to describe anything, even though some such sentences may be used on other occasions to do so.

Now, Boghossian supports his argument for dividing analyticity into vacuous and epistemic variants by attempting to show that a sentence may be epistemically analytic without its truth either having been established by convention or failing to correspond to some worldly fact. To show that instances of epistemic analyticity have neither of those consequences, in both cases, Boghossian cites Kripke's reading of the standard metre case.

As I have shown, Kripke's reading of this argument effectively presupposes at the outset that a norm ("S is a metre long at t_0") must still describe how things are; indeed, it is unclear from Kripke's text whether he recognises that the stipulative definition is supposed to be normative.[86] Recall, in Kripke's scenario, the stick S simply serves as an example of a length which the phrase "one metre" picks out across all possible worlds—that is, in Kripke's terms, it rigidly designates (1972, 55). So "one metre" is a rigid designator, yet the sentence used to ordain the stick with this status, "S is a metre long at t_0", is held to be a flaccid designator, meaning that what it "picks out" could easily have been different: the stick could have shrunk or expanded, e.g. (putting aside, as I am not sure Kripke is entitled to do, the fact that variables such as temperature were historically

---

[85] I do not think this position could only be cast in terms of a correspondence theory of truth, but the details are vague, and this kind of perspective on truth fulfils the purpose for the example's sake.

[86] Note, this is hardly an obvious form for a norm to take. Indeed, it is quite telling that Kripke imagines this form of words for the statement uttered at the crowning moment of the metre standard's introduction, rather than a statement such as "This stick is a metre long".

allowed for in the determining of standards). On Kripke's view, then, it is a contingent matter of fact that the definition "S is a metre long at t_0" is true: the stick S could have been longer or shorter than the length it is supposed to fix the reference of.

But—and this is the crucial point—it is only intelligible that the stipulative definition was a contingent statement if one supposes that it functioned as a description. For descriptions may be wrong; that to which a definite description refers, e.g., could vary across different possible worlds, and so is not "rigid", in Kripke's terms. Accordingly, were the world and its contents arranged differently, any given description (true in "this world") is susceptible of becoming false. However, norms are not descriptions. As we have seen in this chapter, they are prescriptions for how to describe. Consequently, introducing a standard of measurement by the phrase "This stick is (to be called) one metre long" is not to describe the length of a stick (indeed, not to describe such a stick in terms of a unit not yet introduced!); rather, it is to say something akin to "This stick is what we will use henceforth to determine whether some other item is one metre in length."

Accordingly, I cannot find any support for Boghossian's arguments against the metaphysically vacuous aspects of analyticity. An appeal to Kripke's reading of the standard metre case is of no help, since it begs the very question at issue. The question at issue is whether a sentence's being a definition has the consequences either that it is not "in fact" true, or is conventional in nature. Boghossian argues that a definition has neither of these consequences, and does so by citing Kripke, who himself presupposes that something's being a definition does not preclude its being descriptive. Kripke is right to point out that a sentence's being descriptive leaves it susceptible of contingency, and of truth or falsity in fact. But he is wrong to treat a definition as a description, and Boghossian is likewise wrong to reproduce this error.

Moreover, one does not get out of the struggle here by pointing to the pliant and flexible nature of norms. It is true that in practice, formerly established norms have a fading and varying authority over the aptness and correctness of what we say and do. And it is instructive to recognise this. A proper appreciation of how norms actually function helps us overcome simplistic models, for example, of what a word means; how meaning changes; and how disagreement surfaces and is resolved. But none of this takes away from the fact that where a statement is normative, it is not descriptive, and so sentences which express norms—the analytic ones—express in a sense something conceptually true: they express the normative employment of concepts in practice, as the practice stands.

Chapter Three
The context-sensitivity of analytic truth

Like Boghossian, Williamson also holds that analytic truths can be conceived in two main ways—metaphysically and epistemologically—although in his view neither way of conceiving analytic truth makes it any more useful in philosophy. In Chapter Two, I criticised this distinction. For one thing, I pointed out that the distinction is arbitrary; that is, historically, conceptions of analyticity have not fallen either side of this dividing line. For another, I showed how looking at historical conceptions with this distinction in mind only obscures how those conceptions were employed. Finally, and most importantly, I argued that the distinction, on its own grounds, does not hold water. That is because Boghossian's arguments against the so-called metaphysical conception of analyticity (which I prefer to call the vacuousness or non-factual conception), which motivate the distinction in the first place, rely problematically on an untenable presupposition: that sentences used to prescribe may also, at once, describe.

Given that prescription and description are exclusive, I take it that the work of my previous two chapters culminates in the following thought: it should not be controversial to claim that some sentences may be conceptually true; i.e., that they are expressions of practical norms. Accordingly, on my view, much of the literature is simply confused when it talks about "metaphysical analyticity". I prefer to speak instead in terms of the metaphysical *aspect* (or vacuous aspect) of analytic sentences. The supposed metaphysical conception, as we have seen, is often articulated in a worry such as this: How can any sentence be true solely on account of what its words mean? But by emphasising that an analytic sentence simply expresses a norm, I try to diffuse the mystery that inspires and compels this question.

But all the same, there are other, related, concerns—we will see in this chapter specifically that the worry over the purported metaphysical queerness of analyticity resurfaces in Williamson's own attack on analytic truth, in both of its putative general forms (vacuous and epistemological). And so a predominant concern of this chapter concerns what Timothy Williamson has to say about analyticity. In his book *The Philosophy of Philosophy* (2007), Williamson takes aim at both the vacuous and epistemologically distinct aspects of analytic sentences. This attack forms the crux of his negative arguments against any view which regards philosophy as distinctively conceptual. Accordingly, some background is needed first in order to place his arguments in a wider narrative.

In his philosophical methodology, Williamson defends both the claim that philosophy is an "armchair" discipline—i.e., that it does not fruitfully engage in experiment, measurement, and observation (2007, 1)—and that, in spite of this status, its methodology is not to be sharply

distinguished from those of other modes of enquiry. "Philosophical exceptionalism" is the name which Williamson gives to the opposite position, which claims that philosophy's methodology is sharply distinct from other disciplines, like those of the natural sciences; it is this position which he takes as his primary target. Williamson understands exceptionalism as arguing, for example, against naturalism that philosophy does not ask questions about the world (as a posteriori disciplines like the natural sciences do), but rather that it investigates what we bring *to* such a posteriori forms of enquiry; i.e., "our conceptual or linguistic competence" (2007, 2). So-called exceptionalist philosophers are said by Williamson to have taken the "linguistic" or "conceptual turn", depending on whether they take language or thought as the more "fundamental level of analysis" (where thought is, presumably, considered notionally distinct from language), which both usher in the same conclusion: "philosophical questions turn out to be in some sense conceptual questions" (2007, 2).

Now, §3 of this chapter onwards explains and criticises Williamson's arguments against any attempt to say that some sentences are not made true by the facts, or that one could grasp the truth of a sentence as soon as one grasps what it expresses. Williamson thinks that analyticity must prop up any defence of philosophy as an especially conceptual discipline, and so his arguments are naturally of interest to this thesis. However, the earlier sections of this chapter—§§1-2—criticise the general outline and motivation of Williamson's project. This is also directly relevant to my overarching aim because, as suggested above, the thrust of Williamson's work on philosophical methodology is to show that philosophy need have no special concern with concepts, whether that concern is cashed out in terms of subject matter or method.

In spelling out his project, Williamson groups together targets according as they each take the linguistic or conceptual turn; as an antidote to this move, he offers an insight from David Wiggins. However, I think Williamson's use of Wiggins here muddies the waters considerably. Accordingly, in §1.2, I discuss Wiggins's point in greater detail, and show how it is consonant with the methodologies of philosophers who also seem to have taken the turn by Williamson's own criterion. Following this, in §2, I tackle Williamson's arguments against the linguistic and conceptual turns in much more detail; this gives us good reason to think that philosophy really is, contra Williamson, in some sense a conceptual discipline.

More work must be done, however, since, as indicated, Williamson maintains two further points of interest to us here: first, that the best attempt to articulate what "linguistic" or "conceptual philosophers" take themselves to be doing is focussing on analytic truths; and second, that no conception of analyticity can provide philosophy with its sought-after conceptual credentials. Indeed, this is quite pertinent, since, as will be clear by now, I am to happy to say that philosophy

is centrally concerned with analytic truths, or (if the class of analytic truths is felt to be smaller than that of norms in general) with expressions of norms; in my own way, I also sing this line, though to quite a different tune. However, as we will see in §3, it becomes apparent that Williamson does not seriously have in view a conception of analyticity which appreciates normativity. For example, even though Williamson criticises "Wittgensteinians" for apparently proceeding as if "true" were equivocal between analytic and synthetic senses (2007, 54), his argument to show that truth is unequivocal is entirely insensitive to the normative character of some truths. Moreover, his accusation that Wittgensteinians think of analytic and synthetic truth as forming two distinct senses finds no footing, especially given an appropriately normative construal of analyticity.

What comes to the fore, then, by §4 is that the conception of analyticity he criticises is nowhere near sufficiently nuanced. This point grows out of the last chapter, where I emphasised that analytic truths, being practical linguistic norms, are more diffuse than is commonly imagined. A natural concomitant to the thought that normative practices wane and cease to be normative is the thought that they change, bend, and develop offshoots. For all that, normativity remains within contextually restricted areas; indeed, it *must* do so, on pain of semantic meaning and understanding becoming impossible. So if we fail to notice the context-sensitivity of analytic truth, the very notion that some sentence could be analytic in an epistemological rendering looks extremely suspect. And this is just what Williamson argues: that it is always possible for a competent speaker to understand a sentence of his language and yet dissent from it, whether that sentence is characterised as analytic (even logically true) or otherwise. Thus, according to Williamson, it is possible for a speaker to disagree that adults are grown-ups while fully understanding the content of that proposition.

Now, in my previous chapter, I argued that analytic propositions are not, in *fact*, true: they are true expressions of norms. This is to say that one who disagrees that adults are grown-ups is in effect denying a norm.[87] On the face of it, this looks to be a typical case of misunderstanding, since one would be contradicting a norm. I explore this idea as a response to Williamson's criticisms, and in the process I discover again that normativity is far more sophisticated than is accounted for by both Williamson and his opponents (such as Boghossian). Finally, in the last part of this section and chapter, §4.3, I appeal to this subtler picture of practical normativity to undermine Williamson's reliance on a key device in his dialectic—that of semantic competence. Accordingly, I show that Williamson supports his arguments with an untenable view of semantic understanding, according to which one can always be in a position to deny norms without betokening any misunderstanding.

---

[87] If, indeed, that specification is correct about some contextually specified area of discourse; this is a condition I explore in later sections.

As indicated, however, first we turn to look at the general shape of Williamson's methodological project.

§1 The centrality of language

Following Richard Rorty (1967), Williamson uses the phrase "the linguistic turn" to characterise "linguistic philosophy". Those who have held that language is "somehow the central theme of philosophy" have, by Williamson's lights, taken such a turn (2007, 10). According to Williamson, the general feeling among analytic philosophers is that this turn has passed its sell-by date. Yet, he explains, part of his aim with the book is to find out how moribund it really is, and whether we should restore it to health (2007, 10).

§1.1 Who has taken the linguistic or conceptual turn?

To get a handle on what the linguistic turn involves, Williamson briefly discusses how some who have taken it considered language to be central. For instance, he cites A. J. Ayer, who at one time held that the "propositions of philosophy are not factual, but linguistic in character . . . they express definitions, or the formal consequences of definitions" (Ayer 1936, 44). Williamson recounts how Ayer's division of investigation into "descriptions of objects" and "definitions of words" is avowedly derived from Hume's between "matters of fact" and "relations of ideas". Williamson seems to endorse this view of Hume, calling the analytic-synthetic distinction a "linguistic analogue" of Hume's (Williamson 2007, 11). This is, of course, the same exegetical move we criticised in my Ch. 1, §1.2.1. Further, and also problematically, it is home to a way of thinking that presumes the analytic-synthetic distinction is easily extracted from specific projects and moved wholly to others, as one might unearth and replant a growing vegetable in a new patch. This presumed exportability can give rise to a false impression that one can appraise analyticity *tout court* without paying specific attention to how a particular conception is employed. (And by "particular," of course, I mean conceptions more specific than the alleged metaphysical and epistemological variants.)

Ayer's claim that philosophical propositions are linguistic and not factual is classified by Williamson as being about the method of philosophy (cf. Ayer 1956, 1). But moving forward in his survey, he portrays Michael Dummett's Fregean linguistic turn as concerned both with the method and the subject matter of philosophy (Williamson 2007, 12–13). For instance, he points out, Dummett holds that philosophy analyses the structure of thought (not to be conceived psychologically), and that to do this, philosophy must analyse language. And Williamson suggests

that when others maintained instead that philosophy should analyse thought "directly", not via language, so the "conceptual turn" was born, where the constituents of thought are called "concepts" (2007, 13).

Yet according to Williamson, "linguistic philosophers" still nonetheless spoke of concepts (understood as "what synonymous expressions had in common"), and in his view, any philosophy which makes intentionality (i.e., the "aboutness of thought and talk") central to philosophy has taken the conceptual turn too (2007, 13). Having sketched what it means to be a conceptual philosopher, Williamson suggests that not all philosophy should be deemed "conceptual". He argues that if we accept a mind-independent reality then we should accept that concepts only take up a small portion of it, and, given this, that it is by no means obvious that philosophy's task is to analyse only that small section (2007, 14).

Williamson thus views philosophy in this regard as similar to other disciplines, such as those to which we (apparently) have no problem assigning a nonconceptual subject matter:

> The practitioners of any discipline have thoughts and communicate them, but they are rarely studying those very thoughts: rather, they are studying what their thoughts are about. Most thoughts are not about thoughts. To make philosophy the study of thought is to insist that philosophers' thoughts should be about thoughts. (Williamson 2007, 17–18)

As Williamson sees it, however, much contemporary philosophy seems to go beyond "thoughts about thoughts". He maintains that naturalists in the philosophy of mind study thought as part of the natural world, so that those who theorise about qualia have interests "primarily in the nature of the sensations or qualia themselves, not in our concepts of them" (2007, 18). *A fortiori*, he opines, the case holds for contemporary metaphysicians, whom he thinks are concerned to detect what the basic constituents of our reality are and the relations between them—"not to study the structure of our thought about them" (2007, 19). But even here in metaphysical theorising, Williamson notes, "the armchair" remains pivotal, as do our modal concepts; empirical findings may place constraints on the work, but thought experiments nonetheless remain crucial in Williamson's vision of philosophy (2007, 19). According to Williamson, then, the philosopher's questions somehow treat of "the world", yet do so from an armchair. So perceived, philosophy should, he thinks, be considered an "armchair science", of a kind with mathematics, which is "a science if anything is" (2007, 4).

In the course of tracing a short history of the linguistic or conceptual turn, Williamson briefly stops to cite Wiggins. He professes to find in Wiggins an insight that should make us hesitant to pursue "the analysis of thought or language . . . autonomously with any kind of methodological priority" (Williamson 2007, 20). Specifically, Williamson cites the following remark:

> Let us forget once and for all the very idea of some knowledge of language or meaning that is not knowledge of the world itself. (Wiggins 2001, 12)

Williamson interprets this point as saying that in order for us to define, say, natural kind terms, we must be able to point to real-world exemplars. And since, he suggests, "[w]hat there is determines what there is for us to mean", he concludes that when we know what it is that we mean we thereby know "something about what there is" (2007, 20).

This is a bit quick—as exegesis and as an argument. I suggest that if Wiggins's point (as I understand it) is to count against the linguistic or conceptual turn, it remains quite unclear who actually has taken the turn as Williamson conceives it, because philosophers including later Wittgenstein and Austin broadly agree with it. More on that shortly. Presently, let us pad out the Wiggins point with a little detail.

At first glance, the isolated Wiggins line, cited above, does appear to fit with such of Williamson's claims as "few philosophical questions are conceptual questions in any distinctive sense" (2007, 3). And it is true that Wiggins relates his own point to Quine's "Two Dogmas", which at least inspires Williamson's later suspicion of analyticity, so there is some overlap here. The Wiggins point features in his *Sameness and Substance Renewed*, which aims to elaborate a "theory of individuation", or singling-out of continuants (2001, 1). The remark occurs in a very brief section anticipating the use of the terms "real" and "nominal" as employed in a later chapter. This distinction has Lockean origins and it is not important for our purposes, so I will treat it roughly. Wiggins thinks of real definitions (or rather, explanations) as those which we could not perform without making some essential reference to the thing to be explained or defined. Think of how we point to something in the world in order to explain the meaning of a predicate—say, "*This* is a hedgehog". By contrast, he describes a nominal explanation as readily performable without any sort of ostensive or demonstrative device. It is typified linguistically by "if and only if" definitions which specify non-circular conditions, such as "x is a human if and only if x is a rational animal".

Wiggins notes two things about real, as opposed to nominal, explanations (put aside the nominal case, since it is not important here). First, our real "elucidations" of predicates ("is a hedgehog") depend for their correctness on facts about some real-world things or specimens

(hedgehogs). I presume what Wiggins means in this case is that when I try to explain the meaning of "hedgehog", my (real) explanation ought to pick out a hedgehog (or, I should think, some suitable representation of one). Second, however, Wiggins maintains that despite this factual condition on the explanation, "real definitions lie within the province of semantics, and of empirical fact. No lifelike elucidation of ordinary discourse could be contrived without them" (2001, 12). The natural reading of this point as it occurs to me is that ostensive explanations play a role in our linguistic practices. Hence, in some philosophical pursuit, we should not sensibly think that talk of language and talk of the world could be principally separated. For example, I cannot hope to learn a first language independently of learning the world about me, because even though much of language might not require us to attend to the world, still a significant proportion of it does. So much seems innocent and sensible.

Another point, which Wiggins may not have in mind, is that to explain some worldly phenomenon, I will have to do so under some or other understanding of my words:[88] e.g., when I say, "This is a hedgehog", perhaps it is sufficient to point to specimen which is no longer living (at the side of the road or in a taxidermist's display). Or maybe I can use a plastic facsimile to explain what a hedgehog is. In short, we need to know what will count as my pointing to a hedgehog, or indeed as explaining what one is via an ostensive gesture. These normative concerns belong, at least, in the realm of language; to explain what I mean I shall have to make use of linguistic abilities, and to do so I shall very likely have to use parts of my environment as samples. And so yes, we should abandon the view that knowledge of language is not knowledge of the world. At the same time, we ought to acknowledge that speaking thus of our concepts is precisely a way of coming into touch with "the world", with things as they actually are. Accordingly, if philosophy wants to understand, say, the mind, knowledge, or the self, investigating and clarifying concepts of these phenomena (and related ones) is something philosophy ought to do.

Indeed, for Wiggins, examining the use of concepts is a non-negotiable starting point for enquiry, even if (as he sometimes suggests), we could "explicate" those concepts through philosophical work:

> [W]e have rudimentary pretheoretical ideas of identity, persistence through change, and the singling out of changeable things. By means of these, we may arrive at a provisional or first explication of what "same" means and of the actual application of this relation-word. So soon as that is achieved, there is a basis from which to scrutinize afresh and then consolidate our logical and participative understanding of the individuative

---

[88] I take this way of putting things from Travis (2008b), who claims (rightly, in my view) that all our sayings undergo this interpretative nuance, which itself is tied to specific occasions on which we speak.

practices that a thinker's grasp of the concepts of substance, sameness and persistence through change makes possible for him. At the end of this second phase, nothing will be recognizable as the philosophical *analysis* of "=". . . . Provided we do not despise the ordinary ideas by which we conduct the untheoretical business of the individuation and reidentification of particulars, we can remind ourselves well enough of what regulates the principled employment of "=". (Wiggins 2001, 2–3)

This approach respects our discourse and practice as both a starting point and a constraint on the investigation. It seems, contra Williamson, precisely to place some kind of methodological priority on investigating language and linguistic practices. For all that, it does not disconnect our concepts from what they are about. Accordingly, Wiggins is committed to the view that philosophy begins with concepts of phenomena, while also maintaining that investigating concepts is reflective of phenomena. But exactly those two commitments are shared by many philosophers who, by Williamson's criterion, have taken the linguistic or conceptual turn (because they consider concepts central to philosophy).

Hacker contends a similar point:

> The idea that a linguistic investigation of the use of "X" and a conceptual investigation of X were not also investigations into the nature of Xs would have struck analytic philosophers of the day as perverse. (Hacker 2009, 338)

Hacker cites passages from Wittgenstein, Austin, and Hart in which they each make similar claims. For example, Austin held that examining linguistic norms was not "*merely*" to look at words, "but also at the realities we use the words to talk about"; primarily, that is, to bring into focus the phenomena in which we are interested (Austin 1961, 130). Such a view of the nature of philosophical investigation, indeed, is not uncommon. *Especially* so among those philosophers of a certain generation most renowned for focussing on linguistic intricacies. For example, in his review of Ryle's republished work, Constantine Sandis summarises Ryle's concern with the concept of mind(s)[89] as one of honing focus:

> What Ryle [demonstrated] is that one cannot separate questions about minds from questions about mental concepts. The argument here is not, as is sometimes imagined, that the question "what is X?" is a linguistic or conceptual one, but, rather, that we need to settle upon an agreed concept of X before we can even begin to ask questions about it. (Sandis 2011, 456)

---

[89] Sandis criticises Ryle for sometimes talking as though there were a unitary concept of mind (Sandis 2011, 456).

But if the sort of claim that Wiggins makes is also fairly attributed to philosophers who very clearly took the linguistic or conceptual turn as Williamson conceives it (such as Ryle, Austin, and Wittgenstein), and if Williamson thinks this claim counts against that turn, then I cannot see whom he is targeting.

Further, and more worryingly, it is unclear whether Williamson has fairly characterised either what taking these turns consists in, or why it would be problematic to do so. Indeed, Sandis's point is particularly apt, since Williamson, in his second chapter, labours precisely over how to find sense in just this thought: that philosophical questions are conceptual questions (2007, chap.2). Thus perhaps, right off the bat, Williamson's framing of his task misleadingly characterises the concern with concepts in philosophy—at least as that concern avowedly manifests among a major contingent of philosophy.

§2 Was Mars always either dry or not dry?

As I have anticipated, the majority of Williamson's critique of "conceptual philosophy" (2007, 14) focuses on the notion of analyticity; he thinks that attempts to articulate how language or concepts are central to philosophy ultimately narrow down to claiming that philosophical propositions are analytic (see, e.g., 2007, 48). Prior to this part of his critique, however, Williamson considers more general ways in which philosophical statements or questions might turn out to be only conceptual or linguistic questions—especially by investigating the claim that philosophical questions are about language or concepts. He argues that even those philosophical questions which most plausibly would be regarded as being about language or concepts in some way are not, on inspection, really so—no more at least than ordinary, non-philosophical questions are about language or concepts.

§2.1 The Mars question

So what, then, makes a question distinctively *philosophical* for Williamson? Unsurprisingly, Williamson has no definitive answer to this question; after all, he is arguing that philosophy is not special as a discipline; that it has no unique subject matter or method. Nonetheless, the question Williamson takes as a test case, "Was Mars always either dry or not dry?", is, he claims, likely philosophical. His main reasons for considering it such are that he and other philosophers are interested in it; that to answer it, we tend to invoke classically philosophical concerns; and that it is fundamentally tied up with vagueness—a philosophical and logical problem (2007, 24). On that last point, consider the logical form of the proposition about which the question asks:

$\forall t \, (Dry(m,t) \lor \neg(Dry(m,t))$

That is, for any time *t*, at that time, Mars was either dry or it was not dry. Since the predicate "dry" has borderline cases, as do "bald", "heap", "orange", and so on, it is prone to generating the sorts of puzzles which concern philosophers. And as we see shortly, Williamson thinks it essential to answering puzzles about vagueness, such as that posed by the "Mars" question (as I will call it), that we arbitrate between different logics.

For the purposes of investigating the Mars question, Williamson supposes that Mars at one time certainly was wet (or not dry), and that it gradually dehydrated to its current dry state; there is not some specific time at which Mars was first dry or not dry. Accordingly (and plausibly), he says, empirical measures would be of no use to solving the question—finding out that a given volume of water was present on Mars at some time and not another, e.g., would not help, though were circumstances such that water had never been on Mars, of course, we would have an answer.

Since the Mars question's constituent words appear in many non-philosophical questions (though not, one should point out, so arranged), Williamson does not suppose it to be philosophical on account of its vocabulary; nonetheless, the foregoing general reasons apparently attest to its being philosophical. Yet, he argues, nothing about the question should lead us to view it as explicitly linguistic or conceptual: the referents of "Mars" and "dry" are worldly phenomena, he suggests (in his view, Mars and dryness, respectively). Likewise, it is not implicitly so: e.g., the Mars question is not equivalent to some explicitly metalinguistic sentence such as the following:

> Is the sentence "Mars was always either dry or not dry" true? (Does it express a truth as used in this context?) (Williamson 2007, 26)

For, Williamson points out, we can convert and insert apparently unphilosophical questions into this metalinguistic frame (such as "Was Mars always either uninhabited or not dry?") and we would not thereby have shown such sentences to be implicitly about language. Or, more strikingly, if we had shown such unphilosophical questions to be in a sense implicitly about language, then we could equally maintain that the metalinguistic question is in some similar sense implicitly *not* about language (2007, 27).

Williamson pursues other strategies in a bid to show that the Mars question is implicitly linguistic or conceptual, but in each case, ultimately, he shows these strategies not to work. Instead, Williamson changes tack; he proceeds to assess competing answers to the question by considering the different logics they spring from: classical, intuitionistic, three-valued, and fuzzy. Each approach to logic recommends a different truth evaluation of sentences involving vague predicates. Williamson himself endorses classical logic, and so thinks the answer to the Mars question is "Yes": the Mars question is a logical truth, a generalisation over instances of the law of

excluded middle, P or not-P, for various times (2007, 31).[90] Nonetheless, as he admits, such an answer is still contentious and does not rule out further debate on the issue, which prompts a discussion about the different kinds of logic.

So Williamson's point in raising and discussing this question is not predominantly to answer it, or even to assess its possible answers. One point is to show us just why "semantic considerations" (2007, 31) are relevant in answering philosophical questions, albeit questions not themselves about language or concepts. Another, more significant, reason seems to be to foreshadow the conclusions of the following chapters. For Williamson emphasises that different thinkers, each competent with the use of vague predicates, have different beliefs about vagueness and therefore about the Mars question. And so he is urging suspicion over the linguistic or conceptual turn's apparent claim that the method to solve a philosophical puzzle is merely to retreat back to the concepts with which one is competent:

> We already speak the language of the original question; we understand those words and how they are put together; we possess the concepts they express; we grasp what is being asked. That semantic knowledge may be necessary if we are to know the answer to the original question. It is not sufficient, for it does not by itself put one in a position to arbitrate between conflicting theories of vagueness. For each of those theories has been endorsed by some competent speakers of English who fully grasp the question. (Williamson 2007, 39–40)

There are presuppositions here which certainly need unpacking. For example, what is the picture of "semantic knowledge" that Williamson employs? He points out that theorists of vagueness are competent English speakers, as though those who favour a view of philosophical method as especially conceptual or linguistic must have in mind that the best English-speaking philosophers are simply the best English-speakers. Williamson's picture of semantic competency and its relevance to philosophical enquiry must be probed further, in the later sections.

§2.2 Investigating the question

For now, I wish to work with this Mars question—"Was Mars always either dry or not dry?"—and see whether some philosophical approaches towards it are not, contra Williamson's claims,

---

[90] Williamson gives a fuller treatment of vagueness in his book of the same name (1994), along with discussions on other logics and alternative approaches to vagueness. Williamson's own view—the "epistemic" view—is that there is indeed a precise point at which, say, a heap is no longer a heap after the removal of some particular grain of sand, but we simply cannot know it. Even so, he thinks in principle one could, apprised of all the facts know it, which has a Leibnizian ring to it.

usefully or centrally concerned with concepts or language. For it should become clear that Williamson's own attempts to gloss the Mars question as conceptual or linguistic, either by being implicitly or explicitly "about" language, hardly exhaust the ways in which we might require conceptual resources to respond to this question.

Let us think, then, about this question for a short time. Was Mars always either dry or not dry? By hypothesis we take it as given that Mars used to be wet (or not dry, to be precise) and that now it is dry, after the water gradually evaporated from the surface. Surely, then, this set-up all but guarantees that on some times the surface of Mars might more aptly be described as being dryish or "sort of wet"—at least, in practice, that is how we tend to describe states in between these extremes. It does not seem to me that there commonly is a clear guideline for how to apply such intermediate terms, or even their apparently more defined extremes. (It is, moreover, unlikely that we need such guidelines except on particular occasions.) Suppose an astronaut is testing such material; when ought he to say, "This is [sort of] wet"? (Let us suppose, in line with the strangeness of the example, that the astronaut is not interested in more precise measures, such as humidity levels.) A precise rule may not exist, though factors would influence the astronaut's decision; indeed, the same factors as those that might lend the question significance in the first place. That is, the reason for his testing Mars materials for moisture will set some constraints on the wider implications of the substance being dry or not dry.[91]

Thus, one thing to attend to, here, is *why* we reach for in-between words, which may not yet even be recognised by an authoritative dictionary. If one employs a novel term like "dryish" (which is probably not all that novel), one extends the language, so to speak. The need to so extend often responds precisely to feeling ill-served by a given pair of opposing predicates, or a predicate and its negation. (Did the cake taste good or bad? It was somewhere in-between, and so "good" and "bad" do not cut it.) And with respect to the Mars question, we are ill-served by a pair which seem to depict the two extremes in between which exist many sludgy states. There is something rather unnatural, then, right at the start, with constraining answers to the question in the way Williamson wants to. For if we answer that on some time Mars was "dryish", Williamson will press on: "Okay, but is this dryish state dry or not dry?" One wants to reply that one says "dryish" because Mars was not clearly dry or not dry at that time.

---

[91] What I am saying here seems compatible with claims by Avner Baz, who argues that "everyday" questions always have a *point*: "The point of the question guides us in answering it, and everyday questions that are competently raised have a point: they are expressive in some suitable way of some particular interest in the case" (Baz 2012a, 117).

But let us grant this artificial restriction not to use "in-between" words. Even so, we should stop to ask: What recommends a decision one way or the other here? Consider occasions on which it is a pressing matter to decide whether or not some vague predicate has more or less definite boundaries. It seems pretty clear to me that such decisions are usually made with respect to overriding concerns, normally salient within and specific to a given context. For example, is 15°C a mild or not mild temperature? The answer depends at least on whether we are discussing, e.g., climate, cooking, or storage conditions for some heat-sensitive material.

The answer to the question whether 15°C is mild or not mild, of course, might well not be forthcoming, even having explicitly specified the theme of discussion. For even individuals acclimatised, say, to British temperatures could find 15°C warm for a spring day and yet cold for a summer day. Or, if that example is felt to involve too subjective an employment of the predicate, we can consider an apparently more objective one. Both atmospheric and core body temperatures may change by very small (and unfelt) amounts with disastrous consequences, and so in those contexts we have need to be very precise about otherwise vague terms such as "warm" and "hot", and, indeed, would probably stop using such terms in those scenarios. Likewise, recipe instructions can be frustratingly opaque when they use terms such as "low heat", signifying quite different temperatures according to what is being cooked, among other variables.

One should like to say, then, that "warm" *simpliciter*, taken to apply externally to any given context, is a predicate without application—insofar as predicates are intelligibly applicable only within some context.[92] I cannot see "dry" as escaping the same conclusion. Of course, there is much discussion on whether or not parts of language have a context-sensitive meaning or content. Many concede that indexical expressions such as "I" and "here" have such content (e.g., Herman Cappelen and Ernest Lepore's (2005) "semantic minimalism"), and others extend that thought to specific terms such as "know" (as Keith DeRose (1992) does). Some of the semantic contextualist school might then classify "dry" and other vague predicates as indexical (Jason Stanley (2003) thinks contextualists do just this, but Diana Raffman (2005) disagrees). However, the view I recommend is different, and shares much with Travis's (2008b) perspective. Travis argues that the circumstance of evaluating some statement plays an essential role in determining what is said by that statement. Accordingly, with all terms and *a fortiori* with apparently vague terms, context is to determine the content of a statement containing them, because, if detailed enough, context can

---

[92] For a parallel thought, Amie Thomasson (2007) argues that predicates have application and co-application conditions, and that uses which meet none, such as the metaphysician's use of "exists", preclude intelligibility.

clarify the specific understanding of some term used. Williamson's example, of course, gives us *no* context, and so we are left clueless as to what will count as "dry" when the Mars question is posed.

Williamson himself allows a limited role for context, since he recognises that "aboutness", a key notion in his discussion, is context-dependent: " 'About' is not a precise term. On the most straightforward interpretation, a sentence in a context is about whatever its constituents refer to in that context" (2007, 26). I will soon air my issues with this claim, but let us run with it here. It does not seem a gigantic stretch to go from saying that sentences are context-dependent for their "aboutness" to saying that the constituents of sentences (words), and what they express (concepts), take on different significances depending on those same contexts. For example, taking Williamson at his word, let us say that the Mars question is (in some sense) about Mars and dryness. Perhaps, then, one can only get an adequate idea of what dryness amounts to, with respect to what the sentence is about, when we are thinking about dryness and Mars, planetary surfaces (indeed, Earth's surfaces), and so on. That sort of reflection will consist, *inter alia*, of recalling, working with, and critiquing norms that govern our contextual employment of the concept "dry" and cognate concepts over a range of discourse, as I have gestured at above. For instance, it is patently clear that the correctly uttered "This chicken is dry" (in one circumstance) employs the predicate "dry" according to different standards than "The surface of Mars is dry" (or more differently, "Summer this year was dry"), and so some particular, context-dependent standards or norms must be in play if we are to evaluate "The surface of Mars is dry".

§2.3 What is the question about?

As we just saw, for Williamson, a sentence is most plausibly about whatever its constituent words refer to in a context. With that idea in mind, he argues that the question whether Mars was always either dry or not dry is not about words or concepts. Now, I do not exactly want to argue for the opposite claim here. It may be true that philosophical questions are about language or concepts, *on some or other understanding of "about"*. Because that is what must first be established when asking what the Mars question is about: we must first know how "about" is being used in the demand made of us. The point of this subsection, then, is to show how in the process of tackling a philosophical question, we are required to conduct an overview of normative conceptual employment (here, that of "about"), while attending carefully to the contextual nuance of concept-application. In turn, this discussion begins to show in practice, contra Williamson's overriding point, that philosophy is a conceptual discipline—even though that point is likely not best captured by saying that philosophical questions or truths are about concepts.

Gilbert Ryle (1933) once distinguished between different "senses" of "about". His distinctions were far from exhaustive, but I think they open up a path for us to take. First, Ryle identifies the "nominative" sense of "about" which the term enjoys when a sentence's grammatical subject denotes its topic; the subject is thus usually a noun or pronoun. The sentence "Jane drives a Ferrari" is "about(n)" Jane. Second, since sentences can also be about any of what their constituent nouns or pronouns denote, there is a substantive (we might say, predicative) sense; thus, the sentence "Jane drives a Ferrari" is also "about(s)" some Ferrari. However, Ryle points out, in a conversation a sentence is joined with others. Among the various sentences of a conversation, he says, is often some expression which they have in common (if not directly, then in paraphrases and synonyms, he suggests). And so sometimes a sentence is properly said to be about something when the relevant expressions feature commonly among the sentences of the conversation. Thus, "Jane drives a Ferrari" is not only about(s) Ferraris and about(n) Jane, but also conversationally about—about(c)—driving, cars, Jane's hobbies, Jane's preferences, Jane's activities, and so on (depending on the central topic of conversation in which this sentence features). Notice, this sense is distinct from Williamson's, in that a sentence (in a conversation) may be about a topic even though it contains no nouns or other expressions which signify it (Ryle 1933, 87).[93]

How do Ryle's distinctions apply to the Mars question? Of course, we could say, in line with Williamson, that it is about(n) or about(s) Mars, or dryness. And, consequently, that the Mars question is not about(n/s) words or concepts, at least not in the way that a metalinguistic question such as "How should I pronounce, 'Mars'?" perhaps is. But we should stop here and re-examine this supposedly simple relational use of "about"—for how do we decide that the Mars question is about Mars and not dryness, or vice versa? Or is it about them both? When would a question properly be about what its subject denotes but not be about what its predicate expresses? And so on. Ryle provides a clue here when he reminds us that a sentence may be about some topic even though its expressions do not signify it, on account of its being part of wider *context*. What I want to show is that (at this stage of his thinking) Ryle's point is extremely understated: in practice, determining what some sentence is about is always a nuanced and contextual process.

Consider the following example:

> Leo eavesdrops on his partner Shelly, as she talks on the telephone. From the other room he can hear her excitedly say certain things—"Looking forward so much to tomorrow", "Can't wait!", "I need time to prepare"—when he interrupts her all of a sudden. "What are you talking about?!" Shocked, since she was only preparing Leo's birthday celebrations, Shelly replies, "A surprise party

---

[93] Distinguished from the above, so-called linguistic senses of "about", Ryle characterises a last, "logical" sense. The rough idea is that sentences are sometimes about(r) a logical subject in the way that a definite description might on Russell's analysis.

for you tomorrow actually, which you've ruined. What did you think I was talking about…?" Realising his error, Leo sheepishly answers, "I'm sorry. What you were saying sounded… clandestine…"

Let us ask ourselves what Shelly was talking about in the snippets of overheard talk. Leo takes Shelly's statements (that she is looking forward to tomorrow, cannot wait, and needs time to prepare) to be about some planned dishonest behaviour. Shelly, however, clarifies in her response that she was talking about his surprise party (a surprise no longer). Even though Leo is mistaken, both seem perfectly licit answers to the question of what Shelly was talking about. And notably, both answers are drawn from the circumstance the speakers find themselves in. But it might seem like we could take another approach. For Shelly said "[I] can't wait!" which, now we reflect on it, indicates she was talking (elliptically) about her excitement—roughly equivalent to, say, "I'm so excited!". And she is also "looking forward so much to tomorrow"; perhaps that is a statement about her anticipation then. Likewise, she "needs time to prepare", so surely by that utterance she is talking about her needs or her future needs. It looks as though we could have it both ways (and be left none the wiser).

But matters are not so simple. Observe what just happened: we changed our view of Shelly's statements from being (generally) about the party to being (generally) about herself only by escaping from the ground-level circumstantial detail. Now at the abstract level, we examine Shelly's statements independently of all circumstance. Is this legitimate; can we determine what her statements were about without considering that circumstance? Surely, the only way we can decide the matter is to clarify how "about" is used in norm-bound contexts. For that is where the concept we started out with has a home; that concept is the one we want to deploy in other philosophical questions, such as asking what the Mars question is about. Thus, we can only ensure that we are employing "about" intelligibly if we have in view its legitimate uses.

So, back to examples. Suppose Leo's fears were nearer to the truth (but a little off), and that Shelly's hushed tones and excited claims occur in a scenario where she is planning to rob a bank. The reason she says she cannot wait is because she is excited about robbing a bank. Plausibly, her "I can't wait" is just about her plans to rob a bank then; at least, if pushed as to what she is talking about, that is her honest answer. But if Leo is, say, perturbed by the fact that not only is Shelly planning to rob a bank, she is *excited* about it, then the situation may change. Perhaps he asks her why she would be excited about that, and not be worried or feel conflicted about the very idea; if she replies that she knows it is illegal and of dubious morality, but that she still cannot wait, then now it looks as though her repeated "Can't wait!" *is* about her excitement.

This is not a process Williamson wants to engage in. He is happy to bring in context only if we construe "about" as some sort of function, which takes words as inputs and spits out context-

determined referents. Faced with these two different examples, then, Williamson would surely insist that Shelly's orthographically similar statements are about exactly the same things, which are a function of the words used in the statements. But if the foregoing is correct, that is not what "about" amounts to; it is more sophisticated and harder to pin down than he assumes.

Accordingly, Williamson would be making two mistakes in this test case: first, he insists on a use of "about" which is not the one necessarily being used in the examples, and second, he implicitly assumes that type sentences are already about things—purely in virtue of having words in them, linked together in a grammatical way. But all the facts of usage point in the other direction: it is only *token* uses of sentences (i.e., after a manner of speaking, statements) which are about anything, and token uses vary according to the nuance of circumstance; that is just how "about" functions in practice and gets its currency. Thus, the significance of any use of "about" is transparent only when we attend to circumstance, as we do, unthinkingly, all the time.

We should now apply this understanding to the Mars question: what is "Was Mars always either dry or not dry?" about? The short answer is: nothing. It is about nothing insofar as it is treated as a type-sentence, the uses of which we are not to consider.[94] The longer answer is that the sentence may be about any number of things in its token uses, but we should like to know the background and the relevant concerns of the contextual employment of such uses; that was why, in §2.2, we tried to put it into practice with the case of an astronaut, and reflected on norm-bound contextual employments of "dry".

But if the foregoing is correct, why would anyone be tempted to say that philosophical questions are about concepts? A reason is forthcoming: because, considered generally, philosophical questions are of a kind such that to address them we first have to clarify what the concepts they involve amount to: *just as we have to do with "about"* if we are to work out what the Mars question is about.

§3 Williamson's blind spot

To present, we have critically explored Williamson's general claim that we should not think language or concepts are central to philosophy, especially not in thinking that its characteristic questions are about language or concepts. We have uncovered at least two problems: first, according to many philosophers, conceptual investigation incorporates understanding the nature of phenomena, and second, that even to begin appraising the idea that philosophical questions are about language or concepts, we have to clarify the concept of (at least) "about". Likewise, when

---

[94] Compare here Moore's (2009, 120) review of Williamson.

responding to the question "Was Mars always either dry or not dry?" we had to be sensitive to how "dry" is employed in a range of nuanced contexts—paying attention also to the overriding interests of those involved. We see, then, that only by concretising some type sentence—i.e., putting it to use in contexts—are we able to interpret what it says and how we ought to answer it, if we ought to do so. This already intimates why it is right to say that philosophy is a conceptual discipline: in the fields, at least, of metaphysics, philosophy of language, epistemology, and philosophy of mind, to resolve some philosophical problem is partly to examine the concepts expressed by characteristic statements of the problem. And to do so specifically by scrutinising the norms that govern concrete uses of expressions.

So much, then, for Williamson's attempt to undermine claims that philosophy is of a conceptual nature by appeal to what its questions are about. But Williamson's attack on this claim is one part of his wider critique of philosophical exceptionalism, or linguistic/conceptual philosophy. We now face the remainder.

§3.1 Analytic and synthetic senses of "true"

Williamson's main argument features two central claims: first, that conceptual truths are just analytic truths, and second, that no notion of analyticity can support a linguistic or conceptual turn in philosophy.

> the idea of analyticity . . . is still active in contemporary philosophy, often under the less provocative guise of "conceptual truth." The terms "analytic" and "conceptual" will henceforth be used interchangeably. (Williamson 2007, 49)

> The overall upshot [of this critique] is that philosophical truths are analytic at most in senses too weak to be of much explanatory value or to justify conceiving contemporary philosophy in terms of a linguistic or conceptual turn. (Williamson 2007, 53)

The first claim is not much argued for, and Williamson's statement of it is quite striking. For example, he asserts that the common philosophical practice of seeking out conceptual truths, connections, analyses, and necessities reduces to one thing: "[many philosophers] present themselves as seeking far more general and less obvious analogues of 'Vixens are female foxes.' " (2007, 48). This implicit equivalence claim seems rash, for one can readily imagine a philosopher who disagrees with it while maintaining that our paradigm philosophical questions, such as "Can we ever be certain that what we perceive is thus-and-so?" or "What does it take for something to be an object?", are indeed conceptual questions; that the route to solving or dissolving the

problems they engender is a conceptual route, which involves, *inter alia*, articulating conceptual truths (and not analytic truths).

But of course, whether or not analytic truth and conceptual truth amount to the same thing is a matter partly decided by what conception of analytic truth one has in mind; if one is to say that "analytic" means "conceptual", then one ought to have in mind a specific conception. And in the first two chapters of this thesis, I argued that there are many non-equivalent conceptions of analyticity. For example, Kant held that a judgement is analytic when and only when its predicate-concept is contained within its subject-concept. In Kant's case, and contra the rationalists, while some analytic judgements are conceptual ("All bodies are extended"), many conceptual truths are not analytic. Indeed, Kant arranges his whole critical project around the synthetic a priori—a class including arithmetical judgements and those such as "Everything contingent has a cause". By forming synthetic a priori judgements, he claims, we do not rely on experience, and yet we still increase our knowledge. As such, according to Kant, these judgements are uniquely positioned: they constitute the conditions for the possibility of experience; they are those without which we could not apply our concepts.

However, unlike Kant, I am happy to merge his synthetic a priori class into my analytic class (without necessarily agreeing as to which propositions fall under which classes). For example, a proposition such as "Every event [everything contingent] has a cause", which articulates a condition for the possibility of experience, plausibly is itself expressive of a norm, such as "If one applies the concept 'event', one should also conceive its cause."[95] Perhaps, then, Williamson's equivalence claim is apt for my conception of analyticity, though *not* for all. For according to that conception, analytic truths are normative: they are not, as uttered, descriptions of how things are, but rather are preparatory to such descriptions, by being expressions of norms. The propositions (or judgements) which Kant would call synthetic a priori are, to my mind, normative, conceptual truths. Accordingly, I see no outright harm in equating conceptual truths with analytic truths, so long as one makes clear that to be an analytic truth is to be normative (which is not a criterion for Kant, or many others). At any rate, analytic truths are not factual—unlike truths about norms,

---

[95] I make no claim here about the permanent reaches of such a norm, across contexts. If, say, it turns out that the most coherent way to interpret the radioactive decay inside a nucleus is to say that the process of decay is not caused externally, and thus that some events have no cause, then the norm we work with either has to be adapted or restricted to different contexts. Of course, it may be the case that we find a suitable way to describe radioactive decay in causal terms anyway (e.g., that unstable nuclei provide the conditions for a causal process).

such as that there is a given norm at some particular time within some particular community, which G. H. van Wright called "norm-propositions" (1963, 106).[96]

With the concession made that analytic truths (understood in the foregoing way) are conceptual truths, the time is ripe to examine Williamson's arguments for the second claim identified above: that no notion of analyticity (as he sees it) can support a linguistic or conceptual turn in philosophy. Williamson attacks both the metaphysical and epistemological conception of analyticity. Against the former conception, and as mentioned earlier, Williamson offers the same argument as Boghossian (and Chisholm before him) to the effect that the truth of any given sentence can be explained by a statement first of what it means, and second that what it means actually obtains (Williamson 2007, 59). In this section, I focus on a separate (preceding) argument of Williamson's which has in common with Boghossian's argument the same assumption as that which undermines it. That is, like Boghossian, Williamson presumes that the descriptive garb worn by a declarative sentence ensures that it at least tries to describe something, such that it is either factually true or false. Accordingly, both philosophers are closed to the idea of a sentence which only *appears* to describe but which instead expresses a norm, without describing anything.

Perhaps there is some clue to the difference between Williamson's and my approach at the outset. What I am describing in normative terms—that a sentence-in-use which expresses a norm thereby does not describe—Williamson conceives as some putatively vacuous property that an analytic sentence, or truth, enjoys.[97] He terms this vacuous property of analytic truths "insubstantial", but he finds himself hesitant to do so. What defenders of analyticity tend to want, Williamson thinks, is a sense that analytic truths somehow require less for their truth than their synthetic counterparts; in his words, that analytic truths "impose no constraint on the world" (2007, 52). The metaphor is supposed to depict how sentences such as "All adults are grown-ups" are true independently of how the world is: they possess this insubstantial or vacuous property. The reason Williamson is hesitant to use this term "insubstantial", even in describing his opponent's position, is that he does not think it even possible for a truth to impose *no* such constraint, and so the contrast between "substantial" and "insubstantial" cannot, he argues, get off the ground (2007, 54).

---

[96] See Schroeder (2009, 102–8) for a set of arguments to the effect that, with a few notable exceptions, Wittgenstein's grammatical propositions may acceptably be called analytic truths, contra Hacker.

[97] Williamson is far more prone to talking about analytic sentences, propositions, and statements as *truths* than I am, and this may reflect his commitment (and my lack of such a commitment) to the thought that when talking about truth we are only talking about one sort of scenario, in which some fact about the arrangement of the world holds.

Now, we must note that in the arguments Williamson puts forward to demonstrate his conviction that no truth could be insubstantial, he has in mind some conception of analyticity approximating mine. For example, he alleges that many logical positivists and Wittgensteinians hold that "true" can be disambiguated into analytic and synthetic senses and that "they have described [analytic truths] as stipulations, implicit definitions . . . , *rules of grammar* and the like" (2007, 54—emphases added). Against any such views, Williamson wants to "remind" us that "true" is in fact unequivocal between these senses. The argument he offers here, which I have in my sights in this section, runs as follows.

To show that "true" is not equivocal, Williamson embeds analytic and synthetic propositions into disquotational truth-schemata, and then examines them:[98]

> (T) *"P"* is true if and only if P.

If "true" really is equivocal, then there must be analytic and synthetic variants of this principle, he claims. Further, in virtue of the "iff", (T) can be read either left-to-right (" 'P' only if P") or right-to left (" 'P' if P"). So, for example, (T) can be disambiguated into separate principles, where "a" stands for analytic, "s" for synthetic, "lr" for left-to-right, and "rl" for the opposite:

> (Talr) *"P"* is analytically true only if P

and

> (Tslr) *"P"* is synthetically true only if P.

However, while both these principles seem to hold, Williamson points out that (Talr) and (Tslr) are jointly equivalent to a schema for "simple truth":

> (Taslr) *"P"* is analytically true or synthetically true only if P.

And this equivalence, he suggests, should make us suspect that there is really no ambiguity in "truth" to begin with—for (Taslr) apparently makes plain that whether P is analytic or synthetic, it is made true the same way.

Williamson maintains there is a similar problem with the "iff" read in the other direction, which yields

> (Tarl) *"P"* is analytically true if P.

---

[98] These examples are found in Williamson, as are the relevant abbreviations (2007, 54–55).

Henceforth, for simplicity, I will only discuss the analytic variant of the truth schema as read in this right-to-left direction. If we fill in "P", we get an instance such as

(S) "Adults are grown-ups" is analytically true if adults are grown-ups.[99]

This is a claim to the effect that it suffices for the proposition "Adults are grown-ups" to be analytically true purely that adults are, indeed, grown-ups. That sufficiency claim should strike us as extremely odd. It is not the sort of thing which one would expect any contemporary defender of analyticity to believe (though a Leibnizian rationalist might, given the subject-in-predicate principle; see my Ch. 1, §1.1)—that it follows from adults being grown-ups that "Adults are grown-ups" is analytic. If that were so, any given true fact could mandate a true analytic proposition. Of course, Williamson is trying to use this oddness to urge suspicion about analytic truth altogether. However, I think that reflecting on what is so fishy about this example reveals the same error with which we are by now somewhat familiar. To appreciate this point, we will need to observe its Kripkean origins.

§3.2 The inherited error

Recall §2.3 of the last chapter—Schroeder points out that the copula "is", in the sentence "This stick is 1 metre long", is ambiguous between "is *as a matter of fact*" and "is *to be called*" (Schroeder 2006, 246). In his dialectic, it is Kripke who exploits this ambiguity, by supposing that Wittgenstein was mistaken in judging the standard metre in Paris to be neither one metre nor not one metre in length. Not coincidentally, Boghossian appeals to Kripke's argument to support his distinction between "metaphysical" and "epistemological" conceptions of analyticity, and that is why we sought to examine the standard metre case in greater detail in the prior chapter. It is also no coincidence that Williamson regards Kripke as having really begun to knock analyticity off its perch, when he purportedly clarified the differences between apriority, necessity, and analyticity. According to Williamson, philosophers (whom he does not mention) previously took "analytic" to do both epistemological and metaphysical work; however, Kripke showed us that while "analytic" implies "a priori" and "necessary", neither of these latter two imply "analytic", and so that the work left to do for "analytic" was looking scarce (Williamson 2007, 51).[100] There are two points to make about Williamson's relationship to Kripke here.

---

[99] Williamson does not use this example, though he does discuss the general principle (Tarl).

[100] See Kripke (1972, 39) for the original section. Kripke also talks about analytic *statements*, rather than "analytic", which makes more sense, since "analytic" as a term does not imply anything.

First, it is very unclear just what philosophical tradition Williamson has in mind, with respect to the putative implications between statements of different kinds (necessary, analytic, and a priori). To take up our previous example again, one wonders how Kant fits into this picture, in Williamson's eyes. For Kant ultimately saw "analytic" as a logical or semantic property of a judgement that distinguishes it from synthetic judgements, as discussed above. In particular, if a judgement is not only synthetic but a priori too, by Kant's criteria, then it plays a number of roles. It plays a transcendental role by setting out a condition for the possibility of experience, an epistemological role by amplifying our knowledge, and a metaphysical role by being necessary. For Kant, then, a judgement's being analytic would indeed imply its being a priori; yet he argued centuries earlier than Kripke that a judgement could be a priori without being analytic. Surely, the epistemological, metaphysical, and transcendental roles Kant envisioned for analytic judgements were greatly outshone by those he envisioned for synthetic (a priori) judgements. Presumably, then, Kant is not part of whichever tradition Williamson has in mind.

Second, and most importantly, Kripke's greatest influence on Williamson in this matter is intrinsically related to the error which Boghossian later commits. Of course, Williamson sees it otherwise. He positions himself as finishing off a job that Quine began (but failed to do properly), and that Kripke set on course (Williamson 2007, 50–52). Indeed, Kripke had argued for the existence of both necessary a posteriori truths and contingent a priori truths (Kripke 1972, 35–38), and so had tried to loosen apparent links between necessity and apriority. But to my way of thinking, Kripke's telling influence on Williamson is not that he motivated him to dismantle analyticity; rather, Kripke passed on the same error we discussed in the previous chapter and above. Williamson fails to appreciate that norms are not descriptions; that any given expression of a norm cannot simultaneously describe anything. And so, without this appreciation, *of course* analytic truth will look suspect when crammed into truth schemata which have impoverished expressive powers—impoverished in exactly this respect of distinguishing descriptions from prescriptions.

But when one does appreciate the distinction between the normative and descriptive, then one reads (S) much more naturally as saying this:

> (S′) "Adults are grown-ups" is analytically true if adults are to be called "grown-ups".

(And we read "are to be" as expressing a normative connection.) Stated thus, (S') is quite unremarkable. Moreover, it would not commit anyone who uttered it to a descriptive claim about how, in fact, adults are.

The more expressive schema makes plain how Williamson has in effect written into his conception of truth that all true sentences are true *of* some worldly condition (and thus that all true sentences are descriptions), which is sufficient for their truth. And ultimately, this is a variation on Kripke's theme, since Kripke conflates two senses of the copula, one descriptive and one normative. Accordingly, Williamson is here begging the question against, e.g., the "Wittgensteinians" who describe analytic truths as "disguised rules of grammar and the like" (2007, 54). The circular nature of Williamson's argument can be brought further into focus by looking again at (S'), an instance of the more expressive truth-schema. (S') has no direct disquotational replacement for Williamson's "P" as it occurs after the conditional in (Tarl). Instead, I have specified a norm about the correct use of the given terms, which is what would suffice for the truth of an analytic proposition.

We can see, then, how a blindness to the separate logical functions of norms and descriptions leads Williamson to beg the question against some of those to whom he directs his argument. Right out the gate, Williamson denies any independent logical function to rules by squeezing them into schemata that are ill-equipped to handle them as the normative expressions that they are. But what of the assumption that Wittgensteinians see "true" as ambiguous between analytic and synthetic senses? Let us finish by looking at this assumption:

> The distinction between analytic truth and synthetic truth does not distinguish different *senses* of "true": analytic and synthetic truths are true in the very same sense of "true." That should be obvious. Nevertheless, it is hard to reconcile with what many logical positivists, Wittgensteinians and others have said about analytic truths. (Williamson 2007, 54)

One is left to connect the dots here, but Williamson assumes, at least, that Wittgensteinians deem analytic propositions to be true in a different sense to how synthetic propositions are true.[101] Yet it need not follow from how I sought to expand Williamson's truth-schema that I take "true" to be equivocal.

For I have really advocated in this chapter so far that it matters what *kind* of sentence to which one attributes truth. For example, the claim that "Adults are grown-ups" is true cannot be

---

[101] At any rate, those Wittgensteinians who are comfortable talking in terms of analytic/synthetic, and many are not.

evaluated from a vantage point so abstract that the function of the sentence—as used in a particular context—is invisible. Once one knows that the sentence is being used, say, to express a norm, then one is aware, practically, what it amounts to when one says it is true. And so the sense or meaning of "true" need not be various, since the kinds of things to which one may ascribe truth are already various.

This is somewhat similar to how "Jessica is intelligent" entails different things said of Jessica aged two and then five; we may take "intelligent" to have the same sense it always does, but what the claim amounts to will vary with the age of the person to whom intelligence is ascribed. Yet, of course, the tests of intelligence are not entirely disparate—something links them together. Two-year-old Jessica is more likely to be an intelligent five-year-old than is, say, a lower-performing two-year-old. Further, a player with "footballing intelligence" may be terrible at mathematics, yet it is appropriate to ascribe a great mathematician and a renowned midfielder both with intelligence epithets. In the case of intelligence, it is not hard to see that what joins together different performances and tendencies to perform, in respect of intelligence, is to do with reaching certain standards, innovating, economising, and a whole host of performing abilities. (Ryle (1949, Ch. 2) has an interesting discussion of intelligent performances.) But what of "truth"—what do analytically and synthetically true sentences have in common?

It is better to ask: What do we wish to portray, commonly, when we say that eligible sentences are true? The answer is best seen by attending to the practical role of "is true". Now, as Baker and Hacker discuss, Wittgenstein pointed out that saying that p is true is to *affirm* that p, and one can just as much affirm a sentence expressing a norm as a sentence describing some empirical matter (Baker and Hacker 2009, 271). In both cases, one hopes to portray something affirmative, though the kind of sentence which one affirms will affect what this affirmation amounts to. If I affirm a normative sentence—explaining etiquette, I say it is true that "port is passed to the left"—I affirm a norm. If, however, I affirm a descriptive sentence—I judge it is true that "Bill passes the port to the right"—I portray some set of circumstances as really having occurred.[102]

Perhaps, given these above considerations, the adverbial qualifications "factually" and "conceptually" true do not help matters, since they may predispose one to think that these phrases depict different kinds of truth, instead of indicating that *what* is claimed to be true is different in the two cases, even though something is common between their truth-ascriptions; namely, affirming.

---

[102] Doubtless there are plenty of ways to flesh out this portrayal—I hope not to be committed to a particular theory of truth by choosing this one.

§4 Justified by virtue of meaning

The bulk of Williamson's attack on conceptual or linguistic philosophy focuses on epistemological conceptions of analyticity. As discussed in more detail in my prior chapter, the supposed conception distinguishes itself from a so-called metaphysical, or metaphysically vacuous alternative. Defenders of the epistemological conception of analyticity, such as Boghossian, contend along with Williamson that analytic sentences are "made" true by the way that the world is; this is their argument against the vacuous conception of analyticity, which, as I have elsewhere pointed out, relies on the view, deep down, that analytic sentences are descriptive in function. However, the epistemological conception remains committed to another tenet: that while analytic sentences are not made true by meaning alone, they are *justified* by meaning alone. Accordingly, on this view, one is entitled to know,[103] say, "Adults are grown-ups" on account of one's awareness of what the individual words and the sentence as a whole mean.

This conception of analyticity is another that Williamson thinks is ultimately unable to do the work that philosophers require of it, such as explaining philosophy's "armchair" methodology (Williamson 2007, 73). Williamson targets those who hold an epistemological conception with respect to some vocabulary, be that ordinary, logical, scientific, or even moral. The figures he mentions include not only Boghossian, but Peacocke, Bealer, Gentzen, Prawitz, Dummett, Martin-Löf, Brandom, Horwich, Grice and Strawson, and Jackson, among others. It is clear, then, that Williamson thinks the epistemological conception is pervasive.

But it is notable that Wittgensteinians, even caricatures of them, are not generally in Williamson's crosshairs here. In my view, this is likely because of the wider appeal of analyticity's epistemological character than its vacuous, non-factual character, the latter of which is more commonly defended today by those of a Wittgensteinian or Carnapian persuasion. The market for the epistemological character is, so to speak, crowded out with lots of buyers; there are those who, like Severin Schroeder (2009), argue that trivial analytic truths, being norms of representation, cannot be understood without knowing that they are true; but there are also those such as Boghossian (1997) who reject that analytic truths are norms at all while still wanting to save an epistemologically special feature of analytic truths. By contrast, the market for the non-factual character is nowhere near as crowded, meaning that those who accept this feature of analyticity are easier to cordon off as, say, Wittgensteinians. The issue is, then, that many contemporary

---

[103] If not "know", then, say, justifiably believe, be certain of, assent to, etc.—the specific epistemic status is not particularly important for my present discussion.

philosophers, especially those working on analyticity, simply have not observed or fully registered the point that to prescribe is not at once to describe.

Indeed, without acknowledging the exclusivity of description and prescription, those who nonetheless want to emphasise the epistemological character of analytic sentences unavoidably find themselves in trouble. For the central tenet of the epistemological conception—that we can know some sentence is true by knowing its meaning—is much harder to support when we have modelled all sentences on descriptions of facts (and not expressions of norms). And as we will see, Williamson exploits precisely this aspect of such accounts, when he commits himself to a particular notion of semantic or conceptual competence which tells against it. For if every true sentence is true on account not only of its meaning but also because it is true of a fact, then one is always in a position to understand the meaning and miss the fact; this, in my view, is the crucial mistake made by those who argue for an explicitly epistemological conception of analyticity.

Accordingly, it should come as no surprise that recognising the exclusivity of descriptions and prescriptions, and (re-)acknowledging the normative character of analytic sentences, provides us with the key to resolving the problems which Williamson identifies with epistemological conceptions of analyticity. We turn now to these problems, before showing how to resolve them.

§4.1 Understanding-assent links

Williamson attacks the epistemological conception by showing the unviability of "understanding-assent" links, which he proposes must lie at the root of such a conception (2007, 73–74). To explain what such a link is, take for example the analytic sentence "Every vixen is a female fox". According to Williamson, epistemological conceptions of analyticity require that if one understands the sentence, one assents to it. Of course, such assent would not have to be overt, and is even conceived in dispositional terms. Prima facie, that does not seem an outlandish suggestion. For example, were I to dissent from this sentence, most competent English-speakers would, presumably, consider me to have betrayed my misunderstanding of the words involved. However, against the epistemological conception, Williamson denies that anyone who sincerely dissents from "Every vixen is a female fox"—or, more drastically, "Every vixen is a vixen"— necessarily thereby manifests linguistic misunderstanding (2007, 85ff).

Given the prima facie viability of understanding-assent links, then, what tells against them? That is, how might someone understand an analytic statement such as "A bachelor is an unmarried male", yet still dissent from it? One is naturally inclined to think that understanding the sentence brings instantaneously a grasp of its truth. But Williamson (2007, 118) notes that native English speakers often do not consider a man a proper bachelor if he is in a long-term relationship, even

though he is unmarried. Or, he suggests, some might reject that a mother who, after giving birth, undergoes sex-change surgery to become a man is really a female parent (to tell against, I presume, the putatively analytic sentence, "All mothers are female parents"). For simplicity, Williamson examines the following logical truth:

(1) Every vixen is a vixen.

Now, if (1) is analytic, then, according to the epistemological conception of analyticity as Williamson understands it, some understanding-assent link ought to hold. For example,

(UA) Necessarily, whoever understands the sentence "Every vixen is a vixen" assents to it. (Williamson 2007, 86)

But Williamson presents two deviant, native English-speakers who both (by hypothesis) understand (1) yet dissent from it. Let us focus on Williamson's character from the first example, Peter, who thinks that (1) presupposes

(2) There is at least one vixen.

That is, "[Peter] takes universal generalisation to be existentially committing" (2007, 86—Peter can also resort to metalinguistic and theoretical statements of his reasons here, if needs be). Additionally, Peter believes that (2) is false, due to outlandish conspiracy theories about (the non-existence of) foxes. Given the combination of his beliefs, Peter dissents from (1).

How ought we to respond to Peter? Williamson first remarks that Peter is most likely wrong about this matter, not least since (1) is a logical truth. According to the epistemological conception of analyticity, he thinks, if (UA) holds, and Peter dissents from (1), then Peter does not understand it. Yet Williamson wonders what, in this case, conversation with Peter would be like; he thinks it would be clear to us that Peter had misunderstood a constituent term or a combination of terms (2007, 88). But against this, Williamson claims that Peter understands the words "vixen", "female", and "fox" just fine. Peter's conspiracy theory, he says, involves no "semantic deviation". For even though Peter thinks a sentence involving "every x" is only true if an x exists, most speakers, Williamson maintains, share in a similar belief. Accordingly, he suggests that Peter is competent with "every" in normal circumstances. Indeed, Williamson continues, Peter publishes articles in English, submitted to top, refereed journals; he just has odd views on a couple of matters. With Peter's linguistic competence apparently not in dispute, then, Williamson thinks that his evident misunderstanding is not semantic but logical in nature (2007, 91). And as such Peter is taken to be a counterexample to (UA).

A lone counterexample cannot hope to undermine epistemological conceptions of analyticity in one fell swoop. Yet Williamson concocts other examples for other understanding-assent links based on the same base recipe: in every case, he purports to show that it is always possible not to assent to a claim while understanding it nonetheless.

§4.2 The inadequacy of understanding-assent links

We are in a good position at this point to enquire further about the adequacy of so-called understanding-assent links. Remember, we are not in the business of defending an exclusively epistemological conception of analyticity—that simply will not do. But for all that, we are invested in showing that grasp of analytic sentences is significantly unlike grasp of synthetic sentences. Accordingly, it is worth our while investigating whether these understanding-assent links are really fit for purpose. To do so, let us take the tired and overworked bachelor example.

(3) All bachelors are unmarried

On Williamson's model of what defenders of epistemic analyticity propose, if (3) is analytic, then the following understanding-assent link holds:

(UA′) Necessarily, whoever understands the sentence "All bachelors are unmarried" assents to it.

As we have seen, Williamson's strategy is to show that one could always consistently deny a sentence such as that embedded in (UA') without thereby manifesting a misunderstanding. Indeed, if many English speakers would, in fact, refuse to assent to the truth of the bachelor sentence because of (say) how long-term relationships and civil partnerships are understood, is Williamson not right about this? I suggest both that Williamson is right to reject understanding-assent links *en masse* but that he is wrong to suggest that such links are all that could lie at the heart of analyticity's epistemology in the first instance.

We are able to make both of these points by use of some lifelike examples. Suppose Elaine is an English-speaker who dissents from the claim that a specific unmarried man is a bachelor. We should want to pose her some questions, such as "What do you mean 'isn't a bachelor'? That man is unmarried."[104] She might reply as follows:

---

[104] This approach is taken in a similar situation by Grice and Strawson (1956, 150–51); Hacker (2009, 345–46) also cites their article to show the inadequacy of understanding-assent links.

> (Elaine): Well, yes, I know that, but he's one half of a couple who are all but married—they are in many respects just like a married couple; it doesn't feel right to call him a bachelor; he doesn't live the "bachelor" lifestyle.

Elaine is here able to tell us a story about how she conceives of bachelors and unmarried men; how "unmarried" and "married" are to some extent elastic terms. Note that Elaine's conception of bachelor is the kind which may be expected to be in play when discussing the marital status of bachelors, unlike, for example, the notion of a "bachelor of arts", to use Schroeder's example (2009, 86). Accordingly, what is evident in my example is exactly the sort of elasticity that might make a defender of analyticity wobble, and which, naturally, Williamson would like to exploit. (As such, it is much harder to maintain that Elaine's rejection of (3) simply "betokens a divergent understanding" than it would be for the rejection in Schroeder's example (2009, 86).)

Now, to resume the example, we can imagine at this point an interlocutor (say) "Jerry" who could respond in agreement:

> (Jerry): Oh, well if that's what you mean here by "bachelor", then I understand why you dissent from the claim.

By contrast, next consider George, someone who also makes a prima facie strange claim about bachelors, and the same interlocutor:

> (Jerry): Don't worry, being a bachelor isn't so bad.
> (George): I suppose. My dad's a bachelor, and he's happy.
> (Jerry): Really? I thought your dad re-married…?
> (George): Oh, he has.
> (Jerry): Then why do you say he's a bachelor?
> (George): Yeah, 'cos he's a bachelor, y'know, since he's my father and he had that divorce, and…
> (Jerry): [exasperated] No…

George's response is prima facie confusing, especially given the apparent role of his "since". We would want to converse some more with George to clear this up, and if he still cannot offer us an appropriate explanation, I maintain we should have to say he did not have a full understanding of what it is to be a bachelor.

But let us pause here. I have suggested we would be satisfied by Elaine's answers and yet probably be unsatisfied by George's. I doubt that this prediction appears arbitrary to the reader, and for good reason: Elaine is able to explain her statement, and George is not—or not well enough. Because more than simply helping us understand why she said what she said, Elaine gives us particular reason to agree with her. For, as Williamson indeed points out, English-speakers do say such things. My claim, atop of Williamson's, is that while it may be significant that speakers use language in this way, it is yet more significant that they *ought* to do so.

However, we might worry whether English-speakers really ought to adjust their usage so that not only unmarried men but uncommitted men too are to be called bachelors; that sounds radical, or at least strange. In fact, the point here is not that we (English-speakers) all ought to adjust our usage so. The point is rather that "Uncommitted men are bachelors" has *some* normative weight. Recall, in the last chapter, that I distinguished between an idealised kind of normativity and a more realistic, diffuse kind of normativity. The distinction is between, on the one hand, a rule whose authority reaches out over long stretches (perhaps all) of related discourse and, on the other, one whose authority fades somewhat and is itself susceptible to revision in light of the groundswell of other normative activity beneath it. It seems to me that the reality of language fits much better with this diffuse kind of normativity than the idealised version.

For language, in its dynamic nature, incorporates neologisms and other creative uses; it very often reflects novel situations, new purposes. Accordingly, even once adamantine norms become over time much more brittle, and in their gradual deterioration branch open into other, so to speak, "nearly-norms". Such nearly-norms may yet in time harden into overarching and fully-authoritative norms, or they may legislate over restricted areas of discourse, much as slang has a tendency to do (think of "wicked", "bad", "cool", "hot", "safe", and so on). To my mind, given my ongoing training in and sensitivity to English as spoken in my culture, Elaine's usage picks up on some real normativity. At the very least, "All bachelors are unmarried men" is no longer that rigid; our conceptions of singledom, masculinity, commitment, relationship structure, and even marriage have undergone much change in recent decades (cf. Hanfling 2000, 95, n. 1). This is change enough to limit the apparently pervasive reach of the original norm. But such change brings some uncertainty and in turn that can introduce more critical moments in conversation; challenges, confusions, disagreements, and so on will then find footing, because which norm is to govern what we say is not always transparent. (And neither, for that matter, is how to follow a norm.)

Now, one worry, with this talk of diffuse and changing normativity, is that somehow "anything goes"; that correctness and incorrectness are at risk of slipping away. But that is not what I mean to suggest. Rather, what does "go" (i.e., is correct) varies according to the different pockets of discourse in question. For example, to return back to our example, George's "since" cannot be trusted; it gestures at an explanation but fails to deliver on that promise. "My dad's a bachelor, since as a father he divorced and remarried" is not good enough as an explanation because there is no convention which makes it intelligible. No norms or nearly-norms are in place to legitimate George's claim, whereas Elaine plausibly does pick out something normative. This contrast illuminates how the flexibility of language is not total; language, as it were, builds on, and

sometimes breaks, its own constraints. But it has constraints to break; it is not totally free, or it would not be language.

Indeed, the normativity of language is key to Wittgenstein's closely related notion of a language game, which is partly so named because games are practices which are bound by rules (see Baker and Hacker 2005, 49–50). How much freedom do we have to break apart from the language's inherent normative structure? Consider a remark sandwiched between §§35-36 of *Philosophical Investigations* (1953), in which Wittgenstein throws an odd question to the reader: "Can I say 'bububu' and mean 'If it doesn't rain, I shall go for a walk'?" As is his way, Wittgenstein's response is not obviously revealing: "It is only in a language that I can mean something by something". A remark from his so-called *Blue Book* (material dictated to his Cambridge students in 1933-34) helps clarify the point: a sign or a sentence "gets its significance from the system of signs, from the language to which it belongs" (Wittgenstein 1958, 5). The conception of language as evinced here, then, is *systematic* (though not in the sense of being subject to axiomatisation). That is, "bububu" participates in a language if sentences containing it are systematically related to other sentences of the language, such that one could infer from a "bububu" sentence other sentences which follow from it or license it, and so forth. Thus, Wittgenstein's response to the self-imposed "bububu" question is to say that it is not entirely up to him what significance the term is to have—we need to know that "bububu" participates in a language; that there are norms which guide its usage. Perhaps there could be, though these are yet to be spelled out and connected to the rest of language.

Wittgenstein concludes his thought: "This shows clearly that the grammar of 'to mean' does not resemble that of the expression 'to imagine' and the like." That is, a word's having a meaning, and relatedly, our meaning something by a word, is unlike how we have an unrestrained creativity in what we imagine. For in the realm of imagination, rules (to the extent that they have any purchase at all) can bend so easily or disappear altogether, as is the habit of dreams; we can, plausibly, imagine that "bububu" has some determined meaning, but we do not ensure it via this imaginative moment. Consequently, one's ability to introduce a term on a whim, expecting it to become part of the language, is somewhat attenuated, though not impossible. But what about a term with pre-existing normative connections, such as "bachelor"? In such a case, the normative connections of "bachelor" provide a rigidity to its usage with which we must cope and respond to, at least as a starting point. In accordance with the diffuseness of norms, the rigidity need not be everlasting, but it provides a necessary friction without which, as the conversation between Jerry and George shows, effective dialogue is impossible.

Let us bring this all together. George's claim that his father is a bachelor, by virtue of being his father, or having once been a father and divorcee, misfires as a legitimate explanation in much the same way that Wittgenstein's interlocutor fails to establish that "bububu" means something about walking when the weather is fine simply by whimsical stipulation. Both claims do not participate in a language, where a language is understood as a systematic normative practice (or, more accurately, a whole set of changing normative and nearly-normative practices). Far from there being some idealised authoritative norm to undergird these claims, there is not yet even a pocket of practice which begins to carve out such proper usages. Elaine, by contrast with George, plausibly does pick out such a pocket.

So where does this leave the notion of an understanding-assent link? As I foreshadowed, in practice normativity is a diffuse matter, and such links cannot hold independently of a given context: what holds in some region of discourse need not hold in another. Being a competent speaker of English does not require familiarity with all the intricate pockets of discourse and how they connect. Just as one may fail to be "down with the kids" by thinking that they mean "awful and cruel" by "wicked", so one may be utterly perplexed by someone like Elaine, who thinks that an unmarried single man is not a bachelor on account of his committed relationship. And yet in spite of this distance in usage between speakers, all may still be generally competent; this is one reason why the phrase "semantically competent" is unhelpful, at least insofar as one hopes to explain away uses so divergent as to constitute what would normally be regarded as semantic misunderstandings (so I would contend).[105]

And this is where these considerations come to a head. We can see that Williamson needs such a notion of semantic competence to support his attack on understanding-assent links; Peter is competent enough with the English language, we are told, and specifically, let us say, with "every" and "vixen". But even if Peter's general English competency is very good, obviously he can still make mistakes (as we all do). More to the point, he may be competent with "every" or "vixen" in some pockets of discourse but not others. What betokens a linguistic misunderstanding is not necessarily widespread error; likewise, we do not eliminate room for confusion simply by carefully handling various pockets of discourse with a term. Some mistakes are both small and significant. It is this appreciation of subtlety that puts my analysis of divergent usage on a surer footing than Williamson's. For example, whereas I have drawn a distinction between Elaine and George, it does not seem that Williamson has the resources to distinguish in the same way. For on the assumption

---

[105] Boghossian also deems Williamson's talk of linguistic incompetence to be "misleading", specifically when discussing refusal to assent to certain introduction and elimination rules for logical constants (Boghossian 2012, 232).

of George's more general linguistic competence, both he and Elaine are relevantly similar to the character Williamson exploits, Peter: a competent English-speaker who denies a putative understanding-assent link. Moreover, if Williamson cannot distinguish between the likes of George and Elaine, it is unclear how he can avoid a fatal scepticism about semantics: how is it determined that any speaker succeeds in meaning anything by what they say, if there is no norm enjoining what forms of language are intelligible? Likewise, how are we to understand Williamson? Surely, if he wants to be understood at all, Williamson presupposes background normativity.

§4.3 A better picture: understanding and following a norm

Williamson suggests that Peter's mistake falls "well within the range of permissible variation for linguistically competent speakers" (2007, 118). He claims further that such speakers use the words "with their normal English meanings, despite their errors" (2007, 118). Thus, he concludes, that while those like Peter (and, presumably, George) "are ignorant of some facts about the normal English meanings of the words, such ignorance is quite compatible with linguistic competence" (2007, 119). And that, I take it, is the sense to be given to this phrase "permissible variation"— Williamson is saying that we must allow room in our notion of linguistic competence for some semantic errors.

To my mind, this analysis gets a trivial point right, but many serious points wrong. On the one hand, it is easy to agree with Williamson that English-speakers should be able to use their terms in subtly different ways and still be regarded as generally competent. Indeed, I have emphasised something similar in the sub-section above, with respect to Elaine's apparent deviation. There are, I suggested, "pockets" of discourse which overlap and diverge, between which norms may not hold, or by the influence of one region upon another, norms may change, and so on.[106] Accordingly, I have suggested not only that speakers like Elaine are generally competent speakers (whatever quite that amounts to), but that they are also competent with respect to particular terms such as "bachelor".

On the other hand, Williamson is in no position to acknowledge the salient difference between Elaine's competent semantic variation and the deviant uses to which George puts "bachelor". For let us consider where Peter fits in with respect to Elaine and George. It is not altogether clear, since much depends on the further conversation we have with Peter. It may turn out that Williamson's analysis is partly right—that Peter has a conspiratorial, wayward view of vixens, and

---

[106] Chapter 5 focuses in greater detail on the relationship between these pockets, or as I use Mark Wilson's term, "patches" of usage.

a commitment to the logical belief that universal generalisation commits one to existential generalisation. We could say even more about Peter, then: given that there are pockets of discourse in which "every" is taken to have such a commitment, his usage latches onto something part-way legitimate—something conventional, at any rate. In the final analysis, it may turn out that he and others are simply wrong on that matter. Maybe we argue this by extending a point Strawson made in his debate with Russell (Strawson 1950, 331–32). Strawson argued, contra Russell, that definite descriptions, such as "The King of France is bald", do not yield disguised existential statements on logical analysis, such as "There is at least one King of France". Rather, he held, it is simply a common presupposition that when one uses a definite description, there is a unique object to which it refers; one presupposes this and does not state it. In a like manner, perhaps, one makes no implicit existential statement in universal quantification, though one does presuppose it. If that is true, the nearly-norms that Peter latches onto are worth displacing. So be it, and it may turn out that we best express Peter's commitment as misunderstanding "every"—much simply depends on what the normative status quo is, or should be.

Thus, Peter can be relevantly like Elaine in my scenario, and nothing yet seems to have undermined Williamson's position; or he can be more similar to George, and then it becomes clear that something about Williamson's position is fishy. For what can Williamson say about George? I am sure of what "we would say": in everyday dealings with George, when his mistaken belief surfaces—that all fathers, even married ones, are bachelors—competent speakers such as Jerry would most likely cast George's error in terms of a misunderstanding about the concept "bachelor". However, this tack is of course expressly denied by Williamson, because even though George does not assent to (UA′), he is more generally linguistically competent. This is problematic. Because there is a real difference between the likes of George and Elaine, and no account of linguistic understanding is accurate without paying due respect to this kind of difference.

Naturally, then, one wants to enquire about and probe into whatever account of linguistic understanding is really at play in Williamson's view of things. Thankfully, he gives us some pointers. According to Williamson, his argument about the failure of understanding-assent links "exemplifies two interlocking themes" (2007, 91). First, that of epistemological holism, the view that "the epistemological status of a belief constitutively depends on its position in the believer's whole system of beliefs" (2007, 91), and so in Quinean fashion allows unorthodoxy on one point to be compensated for by orthodoxy on many others. Second, that of semantic externalism, the view that "the content of a belief constitutively depends on the believer's position in a society of believers" (2007, 91), and thus explains how, even with unorthodox usage of a term, a person may use that term with its same public (orthodox) sense.

Williamson does not motivate these particular theses in any detail, which is just as well, since he does not take them individually or jointly to support his argument:

> neither epistemological holism nor semantic externalism figured as *premises* of the argument. Rather, the argument appealed to features of the relevant systems of belief that make epistemological holism plausible, and to features of our ascription of beliefs that make semantic externalism plausible. (Williamson 2007, 91)

For my part, then, I do not need to examine, critically, the theses of epistemological holism or semantic externalism. Though it may bear fruit, that is unnecessary work. For if Williamson's general structure of argument against the viability of understanding-assent links is problematic, then it is so much the worse for those other theses. Or maybe such theses can be motivated on other grounds; I need not challenge them in detail here.

Accordingly, the matter we really must investigate is what could prevent one from seeing the difference between George and Elaine, and what the consequences of this omission are. As for the upshot—to my mind, one is left with a picture of language-use, understanding, and socio-linguistic practice which is unbearably deprived of genuine features: there is an intricate web of normativity which is entirely missed out. When we take into account the fact that analytic sentences are norms, and moreover that norms often spread diffusely, harden, soften, change, and improve, then the puzzle over how it is that one could "grasp their truth at once" fades away. If, in some pocket of discourse, "All bachelors are unmarried men" correctly expresses a norm, then one's assent to the claim is part and parcel of one's recognition of it as normatively binding. Understanding and assent comfortably present themselves together, in such cases.

Thus, what prevents one from noticing differences like that between George and Elaine is a tendency to model norms on descriptions, or simply to construe all declarative sentences as descriptions. When that tendency is present, one inclines to believe that analytic sentences must be "made" true by some worldly facts (as I have argued extensively against). At that point, there are two paths to take. If one still suspects that analytic sentences can play an important philosophical role—indeed, an epistemological one—then one must make intelligible an ability of speakers to know that some sentences are true (or be justified in regarding them as such) *even without access* to what in a sense "gives" them their truth (this is what Boghossian tries). The purported path of access is most plausibly thought to be the conceptual or linguistic competence which comes with being a member of linguistic society. Yet this is something wholly bizarre; something, indeed, supernatural. For it would seem to follow that one can intuit knowledge about states of affairs without having any experience of them, even testimonial experience. The picture

is that I learn a language and then somehow I know facts about the world. It is no wonder that this picture is heavily criticised.

The other path to take, when one thinks it impossible for analytic sentences to be true without being made true by the world, is to deny any useful or special significance to analytic sentences. And indeed, when it seems as though an unbridgeable chasm has opened up between facts about the world and our linguistic or conceptual understanding, one only has to exploit it to undermine the Boghossian-style picture described above. For example, by crafting these understanding-assent links and then knocking them down in turn, Williamson capitalises on the independence of meaning from fact, and in turn truth. For this divide helps ensure that when, in his example, Peter has strange beliefs about logic, or about fake vixens, he is not prevented from a full grasp of the meaning of sentences involving logical terms or vixen-related terms.

But once one does away with the troublesome assumption uniting both critics and defenders of so-called epistemic analyticity—that norms are really descriptions—then analyticity no longer appears odd, or as suddenly hopeless in vindicating philosophy's credentials as a specifically conceptual discipline. Because, as I urged in the first section of this chapter, philosophers do investigate the meaning of typically problematic terms; this is a non-controversial rendering of the claim Wiggins makes in his preamble, and which Hacker attributes to a whole range of conceptually-inclined philosophers. One cannot get an insight into, say, the mind without reflecting on mental concepts; the two tasks must be taken in concert. And analytic sentences— norms, more generally—are apt contenders for what philosophy hopes to articulate by investigating the meaning of problematic terms and locutions. Finally, the epistemological and the non-factual character of these norms need not be something mysterious; it only seems it must be so when one fails to realise that a norm cannot describe a fact.

And the consequences of adopting, wittingly or otherwise, the view that norms are descriptions are incredibly damaging. For example, it leads Williamson to oppose understanding-assent links across the board, which in turn makes it unclear how speakers can understand each other: their usage patterns are permitted to vary considerably, such that what before looked like definitions which all should agree upon, are now claims which all are entitled to dissent from, without loss of understanding. He labours, briefly, over this issue: "What binds together uses of a word by different agents or at different times into a common practice of using that word with a given meaning?" (2007, 123). And his answer, save the appeal to causation, is strangely reminiscent of Wittgenstein's "family resemblance" passage (Wittgenstein 1953, §65):

> what makes a unity out of diversity? Rarely is the answer to such questions the mutual similarity of the constituents. Almost never is it some invariant feature, shared by all the

constituents and somehow prior to the complex whole itself – an indivisible soul or bare particular. Rather, it is the complex interrelations of the constituents, above all, their causal interrelations. (Williamson 2007, 123)

Williamson applies this broad sensibility to language more specifically. So applied, it leads him to conclude that we need not share a stock of platitudes to share an understanding of a word. The alternative, Williamson claims, "depends on the assumption that uses of a word by different agents or at different times can be bound together into a common practice of using that word with a given meaning only by an invariant core of beliefs" (2007, 123).

Now, aspects of this very brief sketch might appeal to us. As should be clear by now, I agree that no single sentence can be assented to across all contexts. For one thing, what token utterances of a sentence end up saying is partly determined by salient and background features of a context. For another, norms are subject to revision, to deterioration, to limited and changing jurisdictions— the different "pockets" of discourse. And so sentences expressing norms do not fit the billing of an "invariant core of beliefs" either. Yet the big difference here is that nonetheless there are still matters of right and wrong. Learning a language is, *inter alia*, a matter of being trained in normative structures, where one learns that, e.g., *if* the respective propositions are normative, then an unmarried man may be called a bachelor and an adult may be called a grown-up. As one grows, and as language grows too, one encounters changes to such norms, or finds that they had less reach than in the past. But normativity remains, and misunderstanding still occurs precisely when one violates an active norm.

Moreover, I have barely broached how within one small pocket of discourse we may require different normative standards—compare teaching a child the use of "bachelor" on Monday, and having conversations about bachelors with fellow teachers the day after. On Monday, failure to assent to the bachelor sentence may legitimately constitute misunderstanding, and yet by Tuesday the same dissent need not be taken so, for the fellow teachers can tell a story which is normatively rich enough to explain their own possible dissent, in the way in which Elaine did. This demonstrates not that classic analytic propositions are not normative, or that failure to assent to them does not constitute misunderstanding: it simply shows us how we must broaden our picture of the role, authority, and diversity that norms play in our linguistic practices to make sense of understanding between speakers. This norm-sensitivity is absent in Williamson's discussion; he neglects to treat meaning-understanding as a matter directly related to the contextual employment of norm-bound terms. In the process, Williamson loses the right even to describe his characters as linguistically competent and as understanding each other.

Chapter Four
The transcendental role of analytic truth

In the first three chapters, we investigated and scrutinised conceptions of analyticity as they appear in different places. First, we compared historical conceptions of analytic truth, and found analyticity to be both diverse and pivotal to a number of philosophical projects. In the second chapter, we criticised contemporary attempts to divorce the non-factual character of analytic truth from its purported epistemological upshot—our ability to know these truths a priori, independently of knowledge about empirical matters. We found that arguments to separate apart these non-factual and epistemological features fail to take seriously the normative quality of analytic sentences. That is, Chisholm, Boghossian, and others presuppose that an analytic truth (such as "Adults are grown-ups") is a description and not a rule; or, if they do recognise its normative import, they maintain it is still nonetheless true by virtue of describing the world. However, we presented a series of arguments and imagined normative practices to show that no expression of a rule or norm can ever function descriptively in the way some have imagined.

The value of this point became most salient in Chapter 3, when we analysed Williamson's criticisms of analyticity. It is apparent that he, too, fails to appreciate the normative character of analytic or conceptual truth. With a full appreciation, we were able to see how Williamson's attacks on Boghossian's non-factual and epistemological conceptions of analyticity are as misconceived as Boghossian's distinction in the first place. For Williamson is ultimately unable to explain how it is that speakers betray misunderstandings of specific concepts, or how they can extend concept-usage and develop new lanes for concepts to travel in. And this is because any picture of conceptual truth which does not pay diligence to the normative grounding of concept-use will struggle to make intelligible how speakers of a language both understand and misunderstand the content of each other's concepts.

Further, we found problems with what both Boghossian and Williamson agree upon. For example, both hold that all true sentences are true "of" the world, which is something I have forcefully denied. My denial is founded on the simple point that some sentences express the rules that govern our concepts, even though (as I have also insisted) rules bend, fade, and change. As such, those sentences are not "true of" the world (insofar as this way of speaking is apposite); rather, they are expressions of conceptual norms. The resulting positive idea about truth, here, is that not only is the agreement between descriptions and what they describe but also that between normative propositions and the norms they express (latent as they are in practices) aptly termed "true". We can simply say, then, that predicating truth of a sentence is to affirm it, though what

we affirm will vary according at least to the kind of sentence in question. And so construed, the analytic-synthetic distinction then distinguishes between truths which are normative and descriptive respectively (and need not indicate two different senses of "is true", should that be problematic).

Thus, the last two chapters have provided a defence of the notion of analyticity, or conceptual truth, in part by insisting on its normative character. But we ought now to wonder where all this leaves us. In this chapter, we move on to thinkers from a different strain of philosophy than I have discussed so far, represented by Wilfrid Sellars, Donald Davidson, Robert Brandom, and John McDowell. What instigates this change of course is still the analytic-synthetic distinction, however, and what has been said about it as covered before—indeed, it pays to consider why my approach is one not especially accepted by those we have reviewed from the contemporary discussion of analyticity.

The key to unlocking this puzzle is to change vantage point. I have thus far only looked at contemporary and historic conceptions of *analyticity*. It transpires that to understand fully just why analyticity is contested in the way it is today, we have to look at the other side; the side, that is, of syntheticity or empirical description. In §1, we do just this, and find that latent within Quine's original rejection of analyticity was an untenable conception of the relationship between meaning or language on the one hand, and "the facts" or "the world" on the other. The way this relationship is conceived, moreover, is just as influential as Quine's explicit arguments against analyticity.

This faulty picture of the synthetic is so damaging that, seduced by it, we find ourselves unable to articulate how it is that sentences are ever true *of* the world, and so how what we say is ever vulnerable to the world for a verdict on its correctness. Viewed through this lens, normativity is lost once again, this time as manifest in the way we are, epistemologically, at the mercy of the world before us. In §1, we rehearse arguments from Donald Davidson, who most forcefully argues against Quine that his dualistic way of conceiving mind and world is problematic. But it so happens that Davidson's solution is itself insufficient to restore the normative connection between what we say and that which we say it about. To restore that essential link, we must revisit Kant through the eyes of John McDowell in §2.

The most significant finding of §2 is that McDowell's (dis)solution of Davidson's problem makes indispensable use of "rational relations" between our concepts; and the expression of these, it occurs to me, are grammatical, analytic propositions. Thus, we come to a transcendental finding: without analytic truth, synthetic truth—any true description of the world—becomes impossible. This is a positive finding somewhat distinct from (though related to) the negative arguments against Boghossian and Williamson conducted in prior chapters: a proper appreciation of the

normativity of analytic truth is required for us to make sense of how other truths are synthetic, or empirical.

It is not easy to rest content at this point, however, since an unresolved tension between the various views I have advocated is still present, and this is what I explore in §§3 and 4. The problem is this. First, much of what I have said about the relationship between rational, normative propositions (the analytic ones) and the content of the concepts they involve, resembles, from one point of view, the so-called inferentialism of Robert Brandom. Second, however, Brandom's theory falls foul of exactly the same error as Davidson's by unwittingly severing off the ties between the world of which we speak and our speaking of it. Thus, it seems as though central claims which I advocate about the constitution of concepts threaten to undermine my ability to maintain what we must anyway have: an ability to describe the world truly. What has to give, among this inconsistent set of views, is ultimately Brandom's semantic theory of inferentialism, but it has to be dropped in such a way that still does justice to the relationship between norms and meaning I advocate. By implementing McDowell's critique of Brandom, in §4 we finally reach a perspective from which the link between analytic truth and meaning is transparent, and crucially (unlike Brandom) fully world-involving.

§1 The third dogma of empiricism

As we saw in chapter one, Quine famously attacked the analytic-synthetic distinction, characterising it as empiricist dogma. He thought the idea of some sentences made true by virtue of meaning alone, as opposed to being made true by meaning and facts about the world was untenable because it incorporated a problematic notion of meaning, which was ultimately circular. Moreover, it was undermined, he held, by the falsity of another dogmatic tenet—that sentences face the "tribunal of experience" one by one, in an apparently reductionist fashion. The reality, Quine argued, is not that isolated sentences face this tribunal but rather that the whole of language or (properly) science as a body of theory does (1951, 42). For Quine, then, no sentence in isolation is made true; *a fortiori*, no solitary sentence is made true by meaning alone; rather, every sentence is in principle revisable.

The conception of analyticity which Quine attacked—according to which analytic sentences are true "come what may", independent of extra-linguistic reality—is of course part-Carnapian, especially insofar as it incorporates the idea that the truth of analytic (well, "L-true") sentences can be established alone on the basis of the semantical rules of a language (see Carnap 1947, §2), leaving considerations of fact out of the picture. In my Chapters 2 and 3, I defended the idea that some sentences are not "made true" by the world by pointing out that some sentences, namely

analytic ones, have a grammatical, normative purpose; the argument I made was that sentences expressing norms are not descriptions of the world. In so doing, I avoided the Quinean confusion over how some sentences are not true by virtue of the world. (I am not much attached to this "by virtue of" locution and recognise its misleading nature, but it is easy and inconsequential enough to recycle for now.)

However, there is a deeper problem with Quine's attack on analyticity. Quine's way of framing the relationship between, on the one hand, language, theory, or science, and on the other, the world or facts has the unintended consequence of making the world epistemically inaccessible. All this needs unpacking, which I do presently, but one thought to highlight and hold on to throughout this chapter is that we *must* appreciate the normative dimension of analytic truths to find a route out of the dilemma engendered by Quine's foundational picture.

§1.1 Davidson on the scheme-content dualism

Donald Davidson attacked Quine's underlying view of how language and the world must be related (Davidson 1973, 1983); this is a view which he thinks is encouraged (though not entailed) by rejecting the analytic-synthetic distinction (1973, 9). Davidson himself also finds the analytic-synthetic distinction is useless, as "Quine's faithful student" on the matter (1983, 144), yet he thinks Quine ought to have gone further. For when Quine rejected the analytic-synthetic distinction, according to Davidson he did so in a way which left intact the idea that language is made true by the world, where the world is necessarily not conceived in a conceptual manner but is, as it were, set over against language.

Indeed, Quine had written "[i]t is obvious that truth in general depends on both language and extralinguistic fact" (1951, 36), and he did so as a concession to the framing of analytic truth as an exception to the rule, where the "factual component" of truth is simply null. When Quine rejected the possibility of this null factual contribution, effectively he endorsed the view that, keeping facts about meaning the same, the truth of language thereby depends on *extra*linguistic fact. Falling back on this view was always going to be a risk when giving up the conception of analyticity as Quine had framed it:

> If we give up the dualism [of the analytic and synthetic], we abandon the conception of meaning that goes with it, but we do not have to abandon the idea of empirical content: we can hold, if we want, that *all* sentences have empirical content. Empirical content is in turn explained by reference to the facts, the world, experience, sensation, the totality of sensory stimuli, or something similar. (Davidson 1973, 11)

Of course, what is not in view for Davidson here, but which I think should be, is that it is only giving up a *certain conception* of analyticity which tempts us into keeping the idea that sentences "have empirical content" in the problematic sense (explored shortly). For if we set off, alternatively, by recognising that analytic sentences play a normative role, then we are less likely to make the wrong step of thinking that all declarative sentences are made true by the world.

Still, one may consistently appreciate that analytic truths are normative and yet also think that legitimate descriptions of empirical happenings are, by contrast, made true by what they describe (conceived as "the world").[107] So what is wrong with this view?

Davidson describes this way of conceiving the relationship between the language and the world as a dualism of "conceptual scheme" and "empirical content". The sense of "content" here is illuminated by Davidson's referral to it also as "uninterpreted" (Davidson 1973, 9); that is, a bit of the world which has no conceptual structure—thus not a *description* of the world but the world itself. Indeed, Davidson points out that "content" here is a placeholder for typical expressions of reality, as in "the universe, the world, nature", or alternatively, experience: "the passing show, surface irritations, sensory promptings, sense data, the given" (1973, 14). Content in this sense is something yet to be organised or else "matched" by a conceptual scheme; it is, so to speak, "neutral". The problem with this dualism, according to Davidson, is that such apparently neutral empirical content cannot fulfil the role required of it: it is necessarily not the sort of thing which a conceptual scheme could organise or be made to fit, in a way that secures the truth of our empirical beliefs.

Davidson places these concerns within a particular dialectic in his "On the very idea of a conceptual scheme" (1973), the target of which is a kind of conceptual relativist, represented by Thomas Kuhn, Paul Feyerabend, and indeed Quine. (Relativism in this sense maintains that there could be several incompatible—incommensurable—conceptual schemes.) This is not a dialectic I wish to rehearse, since it is not fully relevant to the chapter. Suffice it to say, however, that Davidson argues against this relativism by using a series of arguments which purport to show this: something could not be a language if it were not in principle translatable into a familiar idiom or language (1973, 7–8). Within this story, then, Davidson's way of fleshing out the problem with the scheme-content dualism is to say that any translation of a person's beliefs from one language to another will require that we individuate what those beliefs are about (events, experiences, sensations, the world) according to "familiar"—i.e., idiomatic—principles (1973, 14–15).

---

[107] I think it is fair to say that Carnap, at several points in the progression of his views, was committed to a similar picture.

The upshot in Davidson's terms is thus that nothing ("no *thing*") on the side of uninterpreted content can "make" sentences or theories (on the side of conceptual scheme) true; however, "*[t]hat* experience takes a certain course . . . if we like to talk that way, make[s] sentences and theories true" (1973, 16). Davidson's explanatory use of the "that-clause" here receives fuller explanation elsewhere, and so we now turn to this material, which is most useful for the purpose of seeing what is wrong with Quine's picture of the relationship between language and world.

§1.2 Reasons for belief

In his later "A coherence theory of truth and knowledge" (1983), Davidson points out that "nothing can count as a reason for holding a belief except another belief" (1983, 141). This moves the former claim about beliefs (concerning how they are individuated) from the ontological realm to the epistemic and rational realm (concerning now what supports our beliefs). But the claims have a common root: a subject's beliefs are connected directly not to the unconceptualised world—the "uninterpreted reality" (1973, 20)—but only ever to conceptual articulations of that reality. A useful example in the later article clarifies this claim:

> Suppose we say that sensations themselves, verbalized or not, justify certain beliefs that go beyond what is given in sensation. So, under certain conditions, having the sensation of seeing a green light flashing may justify the belief that a green light is flashing. The problem is to see how the sensation justifies the belief. Of course if someone has the sensation of seeing a green light flashing, it is likely, under certain circumstances, that a green light is flashing. *We* can say this, since we know of his sensation, but *he* can't say it, since we are supposing he is justified without having to depend on believing he has the sensation. Suppose he believed he didn't have the sensation. Would the sensation still justify him in the belief in an objective flashing green light? (Davidson 1983, 142)

This passage nicely brings out the tension implicit in the idea that our beliefs are connected directly to, and are epistemically supported by, nonconceptual contents—sensations, bits of the world, or unarticulated streams of experience. Davidson's character (call him "Mike") cannot cite any such *things* as his reasons for believing that there is a green light flashing before him. Any attempt to do so will see Mike impose a grammatical structure onto supposedly neutral content, thereby bringing it over to the side of language ("…because I have a flashing green sense-datum"). And if, as Davidson invites us to ponder on, Mike claims not to have a sensation of a green light before him, then in normal circumstances his honest avowal that there nonetheless is a green light before him

is mystifying if not contradictory.[108] Mike is normatively committed, we should think, precisely to disavowing such a belief in the absence of believing in such a sensation, even if the sensation is actually present. This highlights the rational redundancy of the sensation itself in the example, in stark contrast to Mike's beliefs, which we feel must have rational credentials.

It is thus that Davidson rejects the idea of any epistemic intermediary—empirical content— between beliefs and what they are about (1983, 144). Whatever they are about, if beliefs are to provide rational support, they must be expressed in normative, rational terms; that is a requirement which sensations, e.g., fail to meet: "The relation between a sensation and a belief cannot be logical, since sensations are not beliefs or other propositional attitudes" (1983, 143). Yet the idea of a *causal* intermediary, Davidson argues, shows us a way out of this problem, for beliefs are in general, he claims, caused by the world, by the deliverances of the senses, and so on. Therefore, Davidson concludes that via this causal mechanism beliefs do have a basis, albeit not an epistemic one. By this argumentative move, supplemented with some others we cannot explore here, Davidson thinks he has found a reason for thinking our beliefs are largely true (the causal mechanism) which is not evidence (as is putatively provided by, say, sense data) for this truth (1983, 146).

However, as will become clear in the following section, Davidson's move to the causal dimension does not allay the creeping doubt engendered by his argument against the dualism between conceptual scheme and empirical content. On the contrary, it ought to leave us bewildered. As McDowell makes plain, there is something right all along (albeit untenable in the way it is typically expressed) in the empiricist thought that our beliefs are not just caused by the world but are rationally supported by it. The trouble is just how to reconcile that with Davidson's sound (Sellarsian) argument that beliefs are not so supported by a nonconceptual world.

To conclude this section, it is apposite to summarise how we got to this juncture. Quine rejects the analytic-synthetic distinction, and Davidson thinks this was the correct move to make; he opines that erasing that distinction "saved philosophy of language as a serious subject" (1983, 145). Crucially, though, he thinks Quine stopped short of rejecting another fatal distinction—that between scheme and content. He argues that thinking of language or meaning as somehow reducing back to unconceptualised content ultimately opens the door to scepticism and relativism about our conceptual scheme(s), and loses objective truth in the bargain. Moreover, Davidson

---

[108] There is something suspect about "normal" here, since we are granting the intelligibility of having a "sensation" of green. But *that being the case*, the expressed belief "I have a green sensation" counts as a reason for claiming there is green in front of one, while its denial should count against the claim about the green colour in the world itself.

argues it is the way of setting up the analytic-synthetic distinction which brings about this whole mess.

Here, I am in complete agreement with Davidson, with one essential caveat: Quine's way of cashing out the idea of true sentences which are not true on account of describing the world is not the only way, and is indeed fundamentally defective. By contrast, if one sees the distinction between the analytic and synthetic as between sentences playing now normative and now descriptive roles, then a solution to the epistemic tangle presents itself—one which allows us to incorporate the intuitive empiricist feeling that our thoughts really are supported by the world in some sense (a sense which has to be cashed out very carefully).

It is in McDowell that we find the resources to make this case, albeit not cast in terms of analyticity. The larger significance of the foregoing and following arguments for this thesis is thus that not only are there negative arguments to support the view that analyticity is essentially normative; there is a positive, transcendental argument too: if analytic sentences did not express linguistic practical norms, then intentionality—the aboutness of our thought—would not be possible.

§2 McDowell on mind and world

§2.1 The myth of the Given

In this section, we find that McDowell responds directly to themes in Sellars, especially as brought out by Sellars's attack on the "Myth of the Given" (Sellars 1956). Very broadly speaking, Sellars identified the Given[109] in a functional way: anything which purports to play its role—and it is alleged many philosophical items through the ages have tried to play it—is thereby an instance of the Given: "sense contents, material objects, universals, propositions, real connections, [and] first principles" (Sellars 1956, §1). Clearly, having such a range of suspects, Sellars wanted to undermine not just any species of givenness, but what he called the whole "framework of givenness".

According to Sellars, the role any species of the Given plays is explicitly epistemological: it purports to provide justification for a chain of beliefs, to which the chain must ultimately be connected at base, while floating free itself of the epistemic obligation to lean on some further support. One species on which Sellars focuses as an initial way to explore the framework more

---

[109] I follow McDowell's practice of capitalising the problematic Given, to allow room for the fact that it is possible—indeed, necessary—for experience to be given in a more innocent sense: conceptualised experience sets a constraint on what we can correctly say about the world. This will become clearer in due course. See McDowell (1994, 9–10).

generally is the sense-data of traditional empiricist accounts of perception. In essence, then, the Given is the epistemological foundationalist's item of faith, and it is in other terms that which Davidson calls "content" or "empirical content". (Indeed, Davidson (2001, xvi) later acknowledged that his argument should have credited Sellars, given its origin.)

McDowell (1994) portrays philosophy's relationship with the Given as something which seemingly it can live neither with nor without, at least insofar as we appreciate the concern which give rises to postulating the Given. (If we do not recognise this concern, we may be what McDowell (1994, 67) terms "bald naturalists", who opt out of this region of philosophy.) And as we saw with Davidson, this concern arises when we reflect on how our empirical claims, such as "The leaves are green", stand related to their subject matter: are they supported by what they depict, conceived as unarticulated contents?

In this respect, a standard overview of epistemology always includes two contrasting positions: foundationalism, according to which our empirical claims are indeed grounded at some base of certainty, and coherentism (of which Davidson is a former advocate), according to which beliefs connected together in a network provide each other with mutual support, though none is ever epistemically grounded on something exterior to the network. Although McDowell targets both these ways of understanding the relationship between our empirical claims and the world, he does recognise in each position an insight worth preserving. Problematically, however, it is these very insights taken together which generate the concern which bothers him. Nevertheless, his method of resolving this apparent dilemma is to find a way to absorb and recognise both insights.

Foundationalists invoke the Given in response to the felt urgent need to ground our knowledge on something not further epistemically suspect, such as a non-conceptual impact of the world—a "bare presence". For example, one might postulate as a form of the Given a red sense-datum, appearing in consciousness, as supporting a basic empirical claim such as "I see red". This is not unlike, e.g., Carnap's *Aufbau* constructional project—with caveats—which we looked at in Chapter 1; it is also the third dogma of empiricism as characterised by Davidson, which we discussed in the prior section. The insight which partly motivates those who postulate the Given is that somehow judgements must be responsible to the world; that, in order for my claim "I see red" or "there is red before me" to be reasonable, it must be said in response to a reason provided by my experience, or by what I am experiencing. This is a thought, explored more below, which McDowell wants to preserve, in some shape.

Those who reject the myth of the Given, however, such as Davidson, do not give weight to that thought, because an insight of their own seems to tell against it. This insight concerns the nature of rationality; it maintains that reasons stand connected to each other in a *sui generis* sphere.

(McDowell (1994, 67) sees "bald naturalism" as characterised also by its denial of precisely this thought.) We already encountered this insight in Davidson, who articulates it as the claim that "nothing can count as a reason for holding a belief except another belief". Thus, he "rejects as unintelligible the request for a ground or source of justification of another ilk" (Davidson 1983, 141). Anything which appeals to experience to provide that justifying role, must, for Davidson, be succumbing to a form of the mythical Given (McDowell 1994, 14)—we saw, indeed, that Davidson lists "experience" among his problematic categories of reality which putatively explain empirical content. Yet as above, McDowell wants to preserve this Sellarsian thought in some form; he wants to maintain that the world itself, as Given, does not enter into rational relations. And seemingly this is in tension with the empiricist insight above.

According to McDowell, the plausibility of the insights of both foundationalists who postulate the Given and coherentists of a Davidsonian stripe lead philosophy to sway between two positions in an "interminable oscillation" (1994, 9). At one end, McDowell finds Davidson trying nevertheless to provide a stopping point for our beliefs. Granted, Davidson says, experience cannot provide a reason for belief; nevertheless, our beliefs are generally veridical, and are caused somehow by our experience, at which point, being beliefs, they can then enter into rational relations. McDowell explains, however, that the problem with Davidson's anchor—as an alternative to the Given—is that the rocks in which it lands not only fail to justify our judgements, they *exculpate* us from responsibility for them. That is, if we relinquish the Given as Davidson does, it seems we must also relinquish our empirical beliefs' answerability to the world, having thus swung to the other end of the oscillating arc. And so, according to McDowell, philosophy continues oscillating between accepting and rejecting the Given.

§2.2 Concessions to empiricism

This problem and its resolution, on McDowell's account, are essentially Kantian in nature. Recall the very first lines of the (second version) *Critique*, which open with a concession to empiricism:

> There can be no doubt that all our knowledge begins with experience. For how should our faculty of knowledge be awakened into action did not objects affecting our senses partly of themselves produce representations, partly arouse the activity of our understanding to compare these representations, and, by combining or separating them, work up the raw material of the sensible impressions into that knowledge of objects which is entitled experience? In the order of time, therefore, we have no knowledge antecedent to experience, and with experience all our knowledge begins. (Kant 1787, B1)

Kant's opening words convey, *inter alia*, that our experience of the world must involve a passive component. This passivity should be uncontroversial. For example, whether or not the wind blows is out of our control, and so long as we stand outside, we are receptive to the elements; we do not choose to feel the wind. There is of course an amount of our own contribution to this perception: we are standing outside; we are of a certain physical (nervous) constitution that allows us to sense the moving air. But no matter how much of ourselves that we contribute to this experience, that we perceive the flow of the wind cannot entirely be "up to us".

The *Critique*'s opening lines above are perfectly compatible with the thought that experience causes beliefs, and so that our knowledge does begin with experience, even if our judgements about the world are not rationally supported by it. Likely, anyone investigating perceptual knowledge would accommodate this basic thought. Indeed, Kant's concession is clearly agreed upon by both Davidson and any subscribers to the mythical Given. But those who do postulate the Given want to build upon the causal, temporal impact of experience, in a way that Davidson cannot—they want for experience to play an epistemically intermediate function, in Davidson's terms. McDowell calls this requirement a "minimal" empiricism—minimal because he thinks we must assent to it, even if we do not become full-fledged empiricists. Think of it as a second concession to empiricism:

> That is what I mean by "a minimal empiricism": the idea that experience must constitute a tribunal, mediating the way our thinking is answerable to how things are, as it must be if we are to make sense of it as thinking at all. (McDowell 1994, xii)[110]

In McDowell's story, then, Kant shows us the first movement in a relationship between judgement and experience, but in order for knowledge actually to sprout from experience, our judgements need to be more than just caused by it. It becomes clear, to a minimal empiricism,[111] that the judgements we make of the world need somehow to be responsible to it: a subject's utterance of "There is a red ball before me" attains a degree of justification because it is answerable to something in her experience.

The notion specifically of being answerable here is important, and it is not happily substituted by "accurate". That is, we do not satisfy a minimal empiricism simply by saying that a person knows that the ball is red in virtue of uttering "the ball is red" when and only when the ball is in

---

[110] From the introduction, added in 1996.

[111] To be clear, this is *not* a position offered in response or opposition to Kant; on the contrary it is a Kantian point, in McDowell's exposition, as expressed in his dictum that "thoughts without content are empty" (in McDowell 1994, 4).

fact red. That may well be important, as a precondition for environmental knowledge, but it does not show that her judgement meets a rational tribunal in what she encounters in the world. We need instead for the subject herself to perceive that the ball is red, and thereby for her to be rationally entitled to the belief. Thus the problem, as McDowell sometimes puts it, is that subjects need to be responsive to reasons *as such*; that subjects need to be able to recognise and grasp the reasons for their beliefs, even if only when taken back from the experience, temporarily, to reflect on their reasons and whether they provide sufficiently for a belief or action (McDowell 2009, 128). Only this possibility can make it intelligible that the subject responds rationally to her environment—that her judgement is vulnerable to the world for its correctness.

The problem we face in trying to meet this minimally empiricist requirement is that Davidson is quite right to insist that only a belief—or what has a conceptual structure common to beliefs and judgements—can be a reason for another belief. Sellars illuminates this point. Our beliefs and claims, he noted, work within a *sui generis* normative sphere: they are the sorts of things for which we are *responsible*; we *ground* them with good reasons or bad reasons; we *ought* to infer one belief on the basis of another, or a set of beliefs; and so on. In Sellars's visual metaphor, our beliefs operate within the "logical space of reasons" (1956, §36), not the "space of nature". The problem is just how to reconcile this insight about reasons with the first: that the world or our experience of it must provide us, who articulate it, with good or bad reasons for what we say about it.

For given the rational relationships which only beliefs and judgements bear towards each other, how is it possible for something extra-doxastic to provide normative support? That would seem to require that "the world" enters into, or is already a part of, the space of reasons. Prima facie, it is hard to conceive how this is even possible. The carnage of a storm, say, is neither reasonable nor unreasonable—it just *is*. One could intelligibly remark that it is unreasonable, but only in some derivative sense of the term; it would be a metaphorical stretch to describe a storm as unreasonable, and it would not imply that the storm acted with bad reasons (only, perhaps, that God did), or without a reason where it had available reasons. Accordingly, that storms and other worldly forces impinge on our senses is not to be sincerely described in normative or rational language, and therefore it remains totally unclear how anything like the Given—a putative non-conceptual item of experience—could play a normative role.

But *pace* the sceptic, given a minimal empiricism, our empirical beliefs are rationally grounded. And this fact, combined with the insights we have expressed so far, provides us with a transcendental anxiety: how on earth is this possible, given that it must be?

McDowell first expresses this anxiety by working initially with a Sellarsian interpretation of Kant. He expands Sellars's metaphor (McDowell 1994, 5) by imagining now a space of concepts, constituted by our network of thoughts. In metaphor, the claim that the Given (once postulated) could not rationally support our empirical beliefs just is the claim that the space of concepts exhausts the space of reasons; that outside the bounds of the rational is outside the bounds of the conceptual; and, for the coherentist, that experience itself is outside those very bounds. This is the insight McDowell attributes to Davidson.

By contrast, since the Given is defined as non-conceptual intake, those who accept the Given into their epistemic picture—foundationalists—would draw the boundary of the space of reasons wider than the space of concepts. Having done so, however, they too would have absolved themselves of responsibility over their empirical beliefs, because, so long as we accept that the logical space of reasons is *sui generis*, any apparent worldly or experiential "reasons", being non-conceptual, are not eligible for entry into that space. (Think again of §1: the sensation of the flashing light before Mike, or even the flashing light itself, is rationally redundant in an explanation of what makes *him* justified.) The anxiety thus forms because it seems as though neither epistemological picture can allow for what must anyhow be the case: that our experience of the world rationally constrains what we have to say about it. In light of the anxiety, there is an air of mystery to this essential constraint.

In Kantian language, McDowell poses this problem as that of how receptivity (belonging to the faculty of sensibility) and spontaneity (belonging to the faculty of the understanding) jointly contribute to experience. Receptivity covers the passive intake of sensory consciousness, while spontaneity, he says, represents the active, free-thinking nature of reason; "[i]n a slogan," McDowell says, "the space of reasons is the space of freedom" (1994, 5). On this way of viewing the problem, the myth of the Given is a response to the seemingly unbounded liberty of our thinking. Those who subscribe to a form of the Given hold that we are not simply free to say what we like about the empirical world, because items of our experience (as Given), to which we are epistemically responsible (they hope), lay waiting at the bottom of our judgements. And those who reject the Given try to interpret causal relations as rationally restricting what we are free to claim about the world. (It is not only Davidson who tries this route, but Brandom too (1995, 253–54), whom we discuss in this connection in §4.) So the challenge seen in this light is how to weigh down our free thinking with our experiential intake.

McDowell's dissolution is best explored by another colourful example. Imagine that a subject, Rose, faces a green wall, and that, of this wall, she is able to say, "This wall is green". It would be

relatively uncontroversial to say that the abilities Rose exercises when she actively makes such a judgement are conceptual in nature, because here she is operating with the concept of green. Moreover, she can collect instances (shades) under a general notion ("green"); she is aware that green is a colour, of different shades; she would not assent to the claim that green is not a colour, or that the wall is green and red all over, given how it looks in optimal viewing conditions. These are the capacities one would expect Rose to draw on when making active judgements. But stopping to ask Rose about these things by getting her to engage those capacities takes her back from the scene, at a reflective distance; that is the sort of distance that Rose does not have when the world simply impinges on her senses in the everyday flow of experience. It might seem surprising, then, that McDowell's way out of the problem presented is to argue that even in her passive perceptual experience, conceptual capacities are operative: "One's conceptual capacities have already been brought into play, in the content's being available to one, before one has any choice in the matter" (1994, 10).

To my mind, McDowell takes the only route out of a troubling philosophical situation, so long as one has already acknowledged the insights laid out above. However, I do think McDowell's perspective is intuitively more plausible when we ward off an unstated objection right at the start. Clearly, it cannot always be the case that perceptual experience is conceptual in this way. Concepts—conceptual abilities—are learned over time (*pace* Kant's a priori categories), and so there must be a stage in a human's life during which she cannot discriminate the items of her experience, or at least can barely so discriminate. This early stage is summarised by William James's depiction of a new-born's stream of experience:

> The baby, assailed by eyes, ears, nose, skin, and entrails at once, feels it all as one great blooming, buzzing confusion; and to the very end of life, our location of all things in one space is due to the fact that the original extents or bignesses of all the sensations which came to our notice at once, coalesced together into one and the same space. (James 1890, 462)

McDowell's argument, taking account of this objection, should be that as *mature* human creatures we have a rich, conceptualised experience that goes far beyond the relatively indiscriminate "blooming, buzzing confusion" which is so characteristic of the immature experience. The argument is thus best understood not as claiming that conceptual capacities are activated "from the off", so to speak; there has to be a role in place for learning and improving those capacities. (In the view of Wilson (2006), this is an extreme understatement; we accommodate much of his "conceptual evolution" talk in my next and final chapter.)

Let us retrace some steps now in order to refine our understanding of McDowell's resolution. We noted, first, that the logical space of concepts exhausts the logical space of reasons; second, we then saw that those who yield to the myth want for the Given to be located (impossibly) both inside the space of reasons and outside the space of concepts. McDowell is fighting against this myth by urging the thought that experience itself sits within this space of concepts; that all experience is conceptual.[112] And one reason why he urges this is because of the epistemological gain: if experience is correctly so placed, it is thus within the space of reasons, and therefore we can demystify how it is that we justify our empirical judgements and beliefs about the world by reference to it—by reference now to what is "given" in a non-problematic sense.

§2.4 The conceptual content of experience

In order to terminate effectively the uneasy oscillation between accepting and rejecting the Given, McDowell must make intelligible the idea that what experience delivers is already conceptual in nature. In service of this task, one might attempt to construe the dichotomy between spontaneity (our freedom of thought) and receptivity (the passive constraint on our thought) in a new fashion: while all experience is conceptual, spontaneity does not stretch so far as to cover the intake of experience. But, as McDowell notes (1994, 11), this position is not defensible. It retains the spontaneity-receptivity dichotomy, and simply sneaks it into the space of concepts, a space supposedly characterised in the first place *by* spontaneity. So this is (albeit instructively) mistaken, because in keeping hold of the dichotomy, we completely lose sight of the desired epistemological upshot—it is precisely the spontaneous character of our thinking which we must somehow limit and ground, according to a minimal empiricism. What is required, then, to understand how our passive intake can nonetheless be conceptual, as McDowell spells out, is to *dissolve* the dichotomy between ourselves as active, spontaneous thinkers and passive, receptive observers.

This is the crucial thought: receptivity and spontaneity work in concert; the former is not even "notionally separable" from the latter (1994, 41). This claim should not be unpacked as the thought that all passive experience involves active conceptualising;[113] it is not that in order to experience the wall as green, Rose must be thinking about the rational relations between judgements employing the concepts of green, colours, shades, looks, walls, spaces, and so forth. Rather, as

---

[112] With respect to my "baby" caveat, we should take this to mean all *mature* experience, where the requisite level of maturity is for now left unexplored.

[113] Effectively, this is the criticism which Dreyfus (2013) alleges of McDowell—if not with passive experience then with unreflective activity; McDowell (2013) in his response, however, is at great pains to argue that this is not the right interpretation of his views.

McDowell explains, the ability to engage in such active thought is a necessary condition for the possibility of (mature) human experience:

> [E]ven if we consider only judgements that register experience itself . . . we must acknowledge that the capacity to use concepts in those judgements is not self-standing; it cannot be in place independently of a capacity to use the same concepts outside that context. . . . Quite generally, the capacities that are drawn on in experience are recognizable as conceptual only against the background of the fact that someone who has them is responsive to rational relations, which link the contents of judgements of experience with other judgeable contents. (McDowell 1994, 11–12)

Discourse involving "green" is intelligible only against a backdrop of colour-related discourse, which again is intelligible only against a backdrop of discourse about how things seem and look, about the sorts of things which are coloured, about optimal viewing conditions, and about perception as revealing the way things are (1994, 12). It is in these discourses, or networks of thought, that spontaneity is exercised; it is thus where Rose can demonstrate that she comprehends what it is for the wall to be green-coloured.

Accordingly, McDowell is arguing that concepts themselves are partly constituted by the rational relations into which they enter (1994, 13). Moreover, the active spontaneity plays a role in further refining and changing possible intakes of experience. That is, the network of rational linkages in which concepts find their place and distinction is not "sacrosanct": we are responsible for reflecting on that network and adjusting it as needs be (1994, 12–13). McDowell is keen to point out that in principle this flexibility still applies to observational concepts, even though in practice it is hard to see how we might need or be able to revise them.[114]

Indeed, it is somewhat surprising at first that these insights apply even to "close" observational concepts like those of colour. For our colour intake appears most susceptible to being treated (in the manner of the Given) as comprised of self-standing, intelligible items of experience which determine the empirical content of the (colour) concepts sitting closest to them in experience. But, as we see above, one could not engage a capacity to perceive colours *as* they are, in passive receptivity, without being able to navigate along the rational pathways between concepts of colour and ones associated to them in active thought. So even the more passive abilities, which might be thought to play into the hands of a foundationalist, must be conceptual in at least this sense.

---

[114] Perhaps Nelson Goodman's puzzles about the putative colour "grue" (Goodman 1954, 59–83) can help us imagine ways that we might revise our colour concepts, were circumstances importantly different. In the following chapter, we discuss Wilson's blend of inventive and historical examples which suggest that some form of "linguistic engineering" is possible (to use Wilson's (2006) language), even with empirically "close" (to use McDowell's) predicates such as "is hard" and "is red".

Rose's ability to judge that the wall is green draws on the concept of green found in the realm of spontaneous thought. That is a realm in which other propositions about green, colours, things, perceptions, and so forth are rationally related to each other. For example, "Green is a colour", "Surfaces which look green in such-and-such conditions *are* green", etc. Rose is obliged to acknowledge and show a grasp of these rational relations, in order for it to be fully plain that she really does wield the concept of green. And, ultimately, this competence is a *sine qua non* for being able to take in *that the wall is green*, rather than whatever is manifest in simply parroting a verbal response, "green", to certain walls.

§2.5 Analyticity as a precondition for syntheticity

Before proceeding with the next section, it is worth pausing to reflect on what McDowell's train of thought achieves here. We began with a philosophical anxiety, part-inherited from Davidson and before him Sellars, about how perceptual knowledge is possible, especially given two apparently contradictory but plausible insights. This was a problem which on its face was epistemological in character. But in the course of reconciling these insights by paying attention to how experience must be structured, McDowell brings us to thoughts about the nature and content of concepts.

The first thought I want to preserve is that even empirical concepts, such as colour concepts, are at least partly constituted by the inferential relations into which they enter. The second is that in order for a subject to demonstrate competence with these concepts, to count as having them, she must be responsive to these inferential relations by exercising the concepts in the *sui generis* logical space of reasons. To articulate those inferential relations is, at least, to express the norms which partly constitute concepts. It is, as I have urged in the last chapter, to express analytic or conceptual truths, which are essentially normative. Thus, though McDowell does not discuss his rational relations explicitly in terms of analytic truths, the way is very much open to conceiving them as such. With that terminological amendment in place, we come to a conclusion that may be *prima facie* striking: our ability to describe the world truly (or falsely) is predicated on the idea that some other sentences are not "made true" by the world.

Yet this of course needs illumination, and seen in the right light, it is a rather quotidian point. Namely, before we are able to describe how things are, we must learn how to use concepts concerning them; moreover, we must master practices even more general than specific concept-application, such as the practice of describing what we see (or hear; or did see; or might see; or would like to see; and so on!), or aspects of what we see. All such normative practices in a sense pre-empt descriptive moves, though this is a misleading way of putting it. Better: descriptive moves

are only intelligible against the background of normative practices enjoining what is to count as a correct or incorrect descriptive move. Such a background, in my view, endows us with an ability to wield concepts, competently, by teaching us normative, grammatical propositions—in short, analytic truths. Thus, put this way, analytic truth is a precondition of synthetic truth.

The germ of this idea is undoubtedly present in McDowell's *Mind and World*. In his afterword, McDowell discusses how Davidson's rejection of the third dogma ought to bring into question more of the Quinean commitments, including both meaning-scepticism and holism about verification. Accordingly, Quine's argument against the analytic-synthetic distinction need not go through. Yet, McDowell maintains, there is a Sellarsian way of reading Quine which sees his argument against the analytic-synthetic distinction as targeting another form of the Given—this time, an "endogenous" one, local to language or meaning, as opposed to an "exogenous" form local to the external world (such as sense-data or experience conceived non-conceptually, in the manner of Davidson and Quine).[115] That is, one could see Quine as arguing against a linguistic form of the Given (which Sellars discusses too, in less detail); such a species of givenness would be something that purported to supply us with certainty having emanated entirely from within the understanding. One could thus label as a form of the Given exactly the idea that some sentences are made true purely by virtue of meaning alone (see McDowell 1994, 136–37, 157–58).

McDowell thinks the Sellarsian Quine could then find himself in pursuit of a conclusion we want to come to: if some sentences are true purely by virtue of meaning alone, this might suggest a hard, dividing line again between mind (or language) and world, or conceptual scheme and empirical content. Disowning that model of analyticity is then a way of combining receptivity and spontaneity in McDowell's terms. We could say that disowning that idea of analyticity is one way of showing that the output of the understanding is world-involving, much as how discarding the exogenous Given shows that the input of the world is conceptual. Ultimately, McDowell finds Quine lacking here; he is unable to make this move, precisely because he fails to notice the myth of the Given seen from the exogenous side (1994, 136). Accordingly, and as Davidson argued, Quine works with a picture of intentionality which sees language and the world as two disparate entities—the third dogma again.

But what if Quine had noticed and explicitly targeted the analytic as an endogenous form of the Given?

---

[115] See McDowell's *Afterword, Part 1* for general discussion of these points. The endogenous/exogenous distinction is first employed in §3 (1994, 131).

> [I]f I am right that Quine's insight is really a glimpse of the unacceptability of the dualism [between mind and world], perhaps we can rehabilitate the idea of statements that are true by virtue of their meaning, without flouting the real insight. (McDowell 1994, 157)

Effectively, then, McDowell concedes that nothing need be wrong, at the outset, with the very notion of analyticity, so long as it does not involve picturing the space of concepts or reasons and the space of nature on separate sides split by a chasm, which would indeed be "flouting the real insight".

"If," McDowell continues, "the notion of a conceptual scheme need not belong to the dualism, meaning can constitute the stuff of schemes in an innocent sense" (1994, 157–58). And here McDowell arrives only inches away from the conclusion I argue for in this chapter:

> The idea of a structure that must be found in any intelligible conceptual scheme need not involve picturing the scheme as one side of a scheme-world dualism. And analytic truths (in an interesting sense, not just definitionally guaranteed truisms such as "A vixen is a female fox") might be just those that delineate such a necessary structure. (McDowell 1994, 158)

A necessary structure of our conceptual scheme is exactly what I have painted, in prior chapters, analytic truths as providing: not just industry-standard bachelor-sentences, but ones such as "Red is closer to orange than it is to yellow", "The chess king moves one square at a time", "The green light means 'go' ", and "No square is round". Moreover, it is my view that such sentences are expressive of exactly the kind of rational relation that McDowell has in mind in the first instance, when he explains the idea of conceptual capacities being actualised in the realm of spontaneity.

Taking all this into account, then, we can see that how one conceptualises the analytic-synthetic distinction reveals much about how one thinks intentionality is to work. In the fatal case, thinking of sentences as divided into those true by virtue of meaning and true by virtue of the world can make two basic phenomena unintelligible. In the first instance, we may misconstrue the normativity of the analytic—i.e., we may fail to realise that some propositions are grammatical, not descriptive, and thereby prescribe the correct context-bound applications of concepts. In the second instance, we may distort the normative connection between synthetic, descriptive sentences and what they describe—we may, that is, fail to account for our ability to represent, see, know, believe, wish, and desire that such-and-such is the case.

§3 Brandom's inferentialism

I have already foreshadowed that central issues remain to be worked out. First, my endorsement of certain claims about concept constitution—that analytic truths partly constitute the content of relevant concepts—sounds very similar to the inferentialist thesis of Brandom. Second, and problematically, McDowell has shown Brandom's inferentialism to be fundamentally defective, precisely because it allows room for the transcendental anxiety discussed last section, and does nothing to allay it. This second point we will overview in greater detail in §4. The current section will instead focus on summarising Brandom's theoretical machinery and assessing how much overlap there really is with the set of views I have thus far advocated.

§3.1 Inferentialism and analytic truth

We can approach Brandom's inferentialism by considering a notion pivotal to his project—that of a "material inference". Sellars coins the term, and introduces it by contrasting it with formal inference. Formal rules of inference prescribe (and proscribe) the transitions one may make from propositions of one form to another. Take the transition from "The ball is red and round" to "The ball is red". First-year undergraduates are routinely introduced to such transitions when learning the rules of natural deduction systems; e.g., they learn the elementary valid argument form of simplification which holds that the conjuncts of a conjunction may be asserted separately. That rule is thus specifiable purely in terms of form. A *material* rule of inference, however, licenses transitions which are specifiable in terms of the "matter", the content, of the terms involved:

> (a) The ball is red.
> (b) Therefore, the ball is coloured.

As Sellars notes, one way to respond to these kinds of inferences is to regard them as enthymemes—arguments in which a premise is implicit or missing (1953, 313). Once the missing premise is explicitly supplied ("If the ball is red, then the ball is coloured"), no longer does it seem as if there is another rule of inference in play, only a standard modus ponens inference, which secures the argument's validity on purely formal grounds.

To say that there are material rules of inference is to say that the above inference is already valid, without the additional premise. Since validity is standardly defined by reference to inference form, this claim would rely, it seems, on some other notion of validity, which perhaps amounts to an inference which is warranted by a material as opposed to a formal rule of inference (cf. 1953, 313). Sellars considers arguments which might show material rules to be of limited or no importance, but he concludes that "[m]aterial rules are as *essential to meaning* (and hence to language

and thought) as formal rules" (1953, 317). In the following quote, Sellars presents the kind of thought which motivates Brandom, who maintains that material proprieties of inference, by licensing the practical moves we can make between claims, fully explain the notion of semantic content:

> [T]here is an important difference between logical, modal and normative predicates, on the one hand, and such predicates as "red" on the other. In the case of the former, it is obvious that their conceptual meaning is entirely constituted by their "logical grammar", that is, by the fact that they are used in accordance with certain syntactical rules. In the case of the latter, this is not obvious—though, as we are about to argue, it is equally true. (Sellars 1953, 334)

The rhetoric here misleads somewhat, for Sellars is arguing that there is no truly important difference between, say, logical and empirical predicates in respect of how their meaning is constituted—it is only that the meaning of the former is more "obviously" constituted by their logical grammar.

The view that, specifically, the meaning of logical or modal predicates is constituted by the rules legislating correct inferences involving them is a species of what today can be called "local inferentialism" (Murzi and Steinberger 2016): i.e., a semantic thesis about the content of some smaller part of the whole of language, according to which the content of a concept within that domain is constituted by the inferential rules for its correct use. Formal logical vocabulary is sometimes thought to provide the best locality for this thesis, for its purported rules are fairly definite, mostly non-controversial, and easy to discover. Consider the rules held by some to constitute the content of "and": the aforementioned simplification rule which allows one to assert only a conjunct from a conjunction and the rule of conjunction which allows one to conjoin two separately asserted propositions. Given these two rules, someone who agreed, first, that it is raining today, and second, that Donald Trump is not fit to be a US President, is further obliged to agree to the proposition "It is raining today and Donald Trump is not fit to be a US President", on pain of being seen as insincere or confused. Specifically, logical inferentialists claim that, if sincere, such a person would be taken to misunderstand "and".

Indeed, post-Boghossian (1997), this is the point of emphasis for many inferentialists, who take themselves to be committed only to epistemic and not metaphysical analyticity. But this is a distinction we scrutinised in the last two chapters. Briefly, it is thought that no sentence can be true purely in virtue of what its constituent words mean (i.e., no sentence is "metaphysically" analytic), but that some sentences are *known* to be true purely in virtue of knowing the meaning of their constituent expressions (and so are "epistemically" analytic). When example sentences

involve the industry standard (the bachelor sentence), this claim might be thought somewhat plausible, for at least there is a purported subject matter for those sentences to be about (bachelors, or marital statuses); to be true *of*. But the logical inferentialist has different go-to sentences, which are far less obviously "about" anything, since the words, or symbols for, "and", "or", "if", and so on do not on the face of it stand for anything.[116] At any rate, logical inferentialists do take the supposed implicit definition of a logical constant to be epistemically analytic. That is a purported fact they take to explain how it is possible to grasp logical truths.

Contemporary literature is much concerned over whether logical inferentialism is plausible. Scepticism comes from, e.g., Williamson (2007, 2012), who, as we have already seen, denies epistemic analyticity tout court. That is, Williamson argues, if semantic competence cannot ever engender knowledge of a sentence's truth, then logical sentences (such as implicit definitions) cannot be known to be true purely by knowing the meaning of their constituent terms. Accordingly, so Williamson contends, it is always in principle possible not to assent to a sentence—logical or otherwise—without failing to understand that sentence. (In addition to my Chapter 3, also see Williamson 2007, chap.4.)

In light of this ongoing debate, it may seem surprising that anyone is a *global* inferentialist, by taking up the second position which Sellars describes above. A global inferentialist, such as Brandom, has to maintain that the content not just of logical concepts but also, for example, of colour predicates is constituted by the rules for their correct use. Some philosophers worry that a conception of the content of any given concept as involving implicit rules requires substantial, or worryingly, exact agreement in order for communication and understanding to be possible (see Fodor and Lepore 2001).

My first concern at present, however, is to draw attention to how what Brandom has said tallies with what Baker and Hacker had to say, in my Ch. 2, §3, about grammatical propositions. Their analysis of necessary propositions revealed them to be normative,[117] and not, unless used in a different way (*qua* different tokens), descriptive. This was a point which I examined both in the particular case (the "metre proposition") and more generally. Moreover, they held that this normativity makes necessary propositions "concept-forming" and "partly constitutive of the meanings of their constituent terms" (Baker and Hacker 2009, 259). Given that, as I argued, analytic truths may too be regarded as normative, since they also provide "inference tickets" (a

---

[116] This was a key insight of Wittgenstein's *Tractatus*, in which he argued that logical connectives stood neither for relations (1921, §5.42) nor (logical) objects (1921, §5.44).

[117] With the exception of logical propositions, which they hold to be "essentially related to norms of reasoning" (Baker and Hacker 2009, 259).

feature of grammatical propositions), the combination of these claims comes to this: analytic truths are partly constitutive of the meanings of their constituent terms. That is a thought about concept-constitution, but of course there is a related thought: that to understand the meaning of a term is partly to articulate the normative truths which show what one may and may not do with that term.

Brandom makes claims about material rules of inference which are very similar to those just rehearsed about necessary propositions (or analytic truths): "The picture being developed is one according to which materially good inferences correspond to the conceptual content of nonlogical expressions" (1994, 102). Here, Brandom signals his commitment to a claim about (nonlogical) concept-constitution. He derives, again, much of the support for it from Sellars, who holds that material rules of inference—or rather, what they express—are indispensable for both scientific and everyday vocabularies (Brandom 1994, 103).

Brandom's own argument for the importance of material principles is explanatory and pragmatist. Consider the inference from "Today is Wednesday" to "Tomorrow is Thursday". Rather than being enthymematic (and thus, at base, formally valid), he holds this inference to be materially valid. Its validity is, for Brandom, a matter of its being implicit in practical proprieties of inference: in the "moves" that language-speakers license and prohibit by the normative attitudes they take towards the "players" who make those moves (cf. 1994, chap.1). (One can, of course, give expression to the material rule by the use of logical vocabulary, e.g., "If today is Wednesday, then tomorrow is Thursday", even though the norm itself remains implicit.) And so he holds that the correctness of these inferential moves at the practical level is what gives the terms their content. He can then appeal to conceptual content as what underwrites inferences which are materially valid albeit formally invalid. So, in the above example, he argues that the contents of the concepts "Today", "Tomorrow", "Wednesday", and "Thursday" are what make the inference correct (1994, 98). Furthermore, he relates the concept-constitution claim to a corresponding epistemological one: "Endorsing these [material] inferences is part of grasping or mastering those concepts" (1994, 98).

So much rounds off some of the salient similarities between Brandom's picture and what I have so far advocated regarding analytic truths. Brandom's inferentialism, however, is considerably more finely-grained and systematic than the perspectives which I have so far endorsed.

§3.2 Inferentialism: a different game?

Brandom locates a specific area of practice as crucial for conceptual content: the game of giving and asking for reasons. I will detail the structure of this game and its terminology in brief, before making it concrete.

Brandom conceives the Sellarsian "game" of giving and asking for reasons in terms of the practical deontic statuses which "players" possess, and which are instituted by means of the normative attitudes they take towards each other. The two core statuses themselves, he says, are "commitment" and "entitlement" to claims or actions, where actions may too be perceived as the premises or conclusions of practical inferences. Within this game, we confer commitments on ourselves with respect to the assertions we make, through what we say or do; we confer entitlements likewise on ourselves in light of our other commitments, and these entitlements are, effectively, assertions appropriate for us to make. Accordingly, Brandom's "basic" form of "discursive commitment" is the assertion (1994, 157). And finally, assertion, Brandom maintains, is essentially related to inference; i.e., to understand the notion of assertion, one must understand the notion of inference (1994, 158).

That is a very brief overview of Brandom's core model for understanding conceptual content on the basis of inference. It is abstract and broad strokes, so let us consider an example to fill in the details:

> (c) Buttons is a cat.
> (d) Buttons is a mammal.

By Brandom's lights, in virtue of the material norm, which we may make explicit as the rule "If x is a cat, x is a mammal", any speaker who asserts (c) is committed to asserting (d); anyone who is entitled to assert (c) is likewise entitled to assert (d); a speaker who asserted (d) without having asserted (c) would not, on the basis of (d) alone, be committed to (c), though other assertions she makes may commit her to it. In light of the material principle, anyone who asserted (c) but refused to assert (d) would have failed to commit herself to a claim that she was instituted as being committed to, by virtue of the assessments of the interpersonal linguistic community of which she is a part. Everything incompatible with the propositional content of (d) is incompatible with the propositional content of (c), and so, in Brandom's terminology, (c) "incompatibility-entails" (d). And with the notion of incompatibility relations finally in place, Brandom can define the "contents" of any given commitment as "the set of commitments that are incompatible with it" (1994, 160). Thus, the content of a commitment to (c) is an indefinitely long set such as {"Buttons is a dog", "Buttons is not an animal", "Buttons is not a mammal", "Buttons is an invertebrate", . . .}.

The further details of the model described in the sub-section above involve many more intricate distinctions of which we will have no particular use in this discussion. However, I wish to draw attention here to a structural feature of Brandom's system. Brandom is very careful that his vocabularies do not intertwine; for example, incompatibility-entailment makes no mention of formal negation (which it easily could) so that the practice it is used to explain—the game of giving and asking for reasons or the "deontic scorekeeping model of discursive practice"—does not presuppose any formal relations. His description of that practice does not presuppose semantic notions either, such as, crucially, reference. This careful construction allows Brandom systematically to build up a semantic picture about conceptual and propositional content from primitives involving only pragmatic notions such as "material inference", "normative status", "commitment", "entitlement", and "incompatibility":

> The institution of [correct and incorrect moves in a game] by practical assessments on the part of the practitioners is the ultimate source of the meanings of the noises and marks they make, and of the other things they do. (Brandom 1994, 159)

This direction of explanation reflects a general reductivist tendency in Brandom. It is not an issue of which he is unaware (or, as such, an *issue* for him); rather, it explicitly shapes how he frames his task. But it is a bone of contention between McDowell and himself. Exploring some of this disagreement is instructive for coming to an understanding of, first, what conceptual content consists in, and, second, the nature of what is conceptually true: grammatical propositions, analytic truths, and material inferences.

Although Brandom is often characterised as coming from the same school of thought as McDowell, and both as together with Sellars,[118] he and McDowell have had several disagreements in print over core features of the other's philosophy. For example, while Brandom has expressed the thought that his own work not only chimes with McDowell's, but in substantive ways continues and repairs it, or expands on its consequences (Brandom 1995, 248), McDowell explicitly denies this characterisation (McDowell 1995, 1998, 2005):

> I reject the suggestion that my short book [i.e., (McDowell 1994)] issues promissory notes redeemed in Brandom's long one [i.e., (Brandom 1994)]. In crediting me with a

---

[118] See Maher (2012).

sketchy version of the sort of philosophical activity he engages in, Brandom ignores the
dialectical organization of my book. (McDowell 1998, 403)

§4 Bringing the world into the picture

The main reason McDowell rejects the apparent affinity between himself and Brandom is that he
fails to afford the representational dimension of experience any significant role. We have already
seen that McDowell argues that this representational dimension is necessary for empirical
conceptual content (that is, for the content of concepts familiarly used to depict phenomena
encountered in experience, such as "red" and "tree"); it is necessary, he argues, for us to make
sense of our taking a stand on how things are in the world. In this section, I show that this is
something which Brandom cannot accommodate by using the notion of inference (and the game
of giving and asking for reasons) as an unexplained explainer. Largely, that is, Brandom cannot
make room for the role of representation in his semantics precisely because he is committed to
privileging inference at the expense of representation. Whatever stance we take, then, on how
concepts are constituted, it cannot reduce purely to some specified set of inferences, conceived
independently of representation.

§4.1 Brandom's reductivism

An important disagreement between Brandom and McDowell centres on what Brandom calls "the
semantic order of explanation". For Brandom, the link between pragmatics and semantics is
intricate. Roughly, semantics concerns contents, especially intentional contents—the contents of
beliefs, desires, and so forth; pragmatics concerns the significance of those contents, as borne by
intentional states, attitudes, and performances (1994, 68) and as gleaned by a study of practices
involving them. Moreover, practice, as something *done*, holds out the hope of grounding meaning
or content (1994, 91), since practice can be described, according to Brandom, without the use of
semantic terms such as "true" or "represents".

When we have in mind, then, how Brandom construes the relation between semantics and
pragmatics, it is easier to see what he means by the "semantic order of explanation" as a dimension
along which his inferentialism can be articulated. The inferentialist order of explanation requires,
first, that he home in on what the practical significances of our beliefs and such really are, to reveal
their contents, and, second, that he explain how other semantic features of language can be
explained in terms of those (inferentially articulated) contents.

Brandom's contrast to the inferentialist order of explanation is the so-called
"representationalist" version. The representationalist tradition, according to Brandom, takes

representation as its primitive or "master" concept (1994, 6), which means that the contents of concepts and propositions are articulated representationally—by a specification of what is represented by those concepts and propositions. For singular terms, that could be the objects in the world designated by them; for predicates, that could be properties of objects; for propositions, that could be truth-values or sets of possible worlds. The details do not matter, here, insofar as the picture is clear: according to the representationalist order of explanation (as Brandom conceives it[119]), concepts acquire their contents through their representative function, while other semantic notions—such as inference—are explained in terms of those representational contents.

In view, then, are two apparent "traditions" which order semantic explanation in opposite directions.[120] Brandom understands the representationalist tradition as being Cartesian in origin, and often designational in method (that is, representation is cashed out in terms of designation relations). Brandom argues against this order of explanation, but praises the inverted, inferentialist order which by contrast takes the notion of inference to be primitive. He claims that the inferentialist tradition responsible for inverting the order arises in pre-Kantian rationalism, grows in Kant's own philosophy, matures in Hegel, lies implicit in Frege, and finally makes itself explicit in Sellars. The details and accuracy of Brandom's rational reconstructions are certainly questionable. (For example, McDowell (2005, 130–33) charges that portraying Frege as having inferentialist credentials is a severe misinterpretation.[121]) Still, Brandom's inferentialism stands or falls on its own merit; in particular, on its commitment to a reductive order of semantic explanation.

However, as foregrounded in §3.3, it is precisely this defining feature of Brandom's inferentialism which McDowell contests: that one must take either representation or inference as primitive and explain the other in terms of it. Brandom holds that either primitive must be taken as antecedently intelligible (e.g., 1994, 94)—that is, as making sense independently of and prior to other notions. But at first glance the matter is slightly obscured, because, as we have seen, Brandom sets up the divide between pragmatics and semantics in a very particular way: in the first chapter of *Making It Explicit*, Brandom settles on a "normative pragmatics", whereby content is conferred

---

[119] McDowell (2005) argues that Brandom's attack on this apparent tradition has straw men as its targets. Indeed, McDowell deems it near-fatal that this tradition may not exist, since without it—and its mistakes— not much else recommends its mirrored inversion.

[120] To anticipate the next chapter, it strikes me that both of Brandom's supposed traditions adhere to the "classical view of concepts," as defined by Wilson (2006); both, it seems, buy into some notion of "semantic finality". See my Ch. 5, §2 for an explanation of these terms.

[121] McDowell (2005) also attacks Brandom's interpretations of the others. As does Michael Kremer (2010).

on concepts and propositions by the normative attitudes and exchanges of individuals in a community. So it does not seem right to say that Brandom conceives the notion of inferential correctness as intelligible antecedently *of anything*—as Brandom accuses the representationalist tradition of so taking representational correctness (1994, 6f)—but actually as intelligible in light of something expounded on in depth.

And so one might say that the semantic order of explanation which Brandom urges us to adopt is not exactly reductive, because it bottoms out not in semantics but in practice; and, one might reasonably conjecture, concepts could not receive their content *except* through practice. (Wilson (2006) would strongly object to this claim, as we see in Chapter 5.) That is, if it is reductivist to claim that conceptual content is derived from practice, we might wonder if that is simply a reduction worth making.

I think that this is, in rough outline, the shape of Brandom's way of thinking here. Indeed, a central remark of his manifests it: "Semantics must answer to pragmatics" (1994, 83).[122] But I also think it would be rash to see this claim as essentially reductivist or inferentialist—in the sense of Brandom's inferentialism. This is because the claim that concepts could not receive their contents except through practice is vague and platitudinous; it can be unpacked differently to suit the needs of many incompatible philosophies. For even a direct reference theorist will conceive practice in some way so as to complement and fill in the gaps of a semantic story; Kripke's emphasis on baptisms and causal-historical chains of use spring to mind (Kripke 1972, chap.3). Yet the platitudinous claim itself does not get Brandom off the hook (if, indeed, being a reductivist is a bad thing), since his unpacking of the claim *is* nevertheless reductivist, as we will see presently.

Recall from §3 that Brandom's conception of discursive practice relies on the normative material principles which underwrite the inferential transitions speakers make. This relation itself is reductively conceived: "material proprieties of inference have been treated as primitives, playing the role of unexplained explainers" (1994, 133). Now, according to Brandom these primitives are instituted by players in the game of giving and asking for reasons, or the "deontic scorekeeping model of discursive practice", where the basic move (or discursive commitment) is making an assertion, and where assertions are understood in turn by the more basic notions of (assertional) commitment and entitlement. And, finally, these notions are spelled out in terms of incompatibility relations between propositions. (For example, again, "Buttons is a cat", in Brandom's terminology,

---

[122] It is, indeed, the section title in which the claim figures. McDowell also focuses on this insight (1997, 157), by showing that reasonably interpreted it ought only to commit Brandom to a position that he himself identifies as possible, "weak inferentialism" (1997, 160); see *Making It Explicit* for the weak/strong/hyper inferentialism distinctions (1994, 131).

incompatibility-entails "Buttons is a mammal".) Brandom claims that the players in this game are able to keep track of each other's commitments and entitlements because of their awareness of the proprieties governing which transitions are allowed and disallowed; thereby, they "keep score" of the discursive field of play. Accordingly, Brandom uses this set of key terms to exhaust the game of giving and asking for reasons, which he takes to model, in principle, real-life social discursive practice (cf. 1994, 158). And this is the problem: discursive practice as we know it could not be conceived in such a set of basic terms.

One huge issue is that it is hard to see how the notion of assertion, which is pivotal to Brandom's deontic scorekeeping practice, is recognisable *as* assertion, given he is only to describe the game in the terms discussed above. McDowell has forcefully argued this point (2005, 126–29):

> To my ear, we have locutions that are explicitly representational as soon as we have "that" clauses, as soon as we have the idea of propositional content. If someone is said to assert that things are thus and so, she is thereby said to represent things as being thus and so. (McDowell 2005, 126)

To understand the force of this point, remember that Brandom by design postpones an account of objective purport, so that representational locutions[123] are prohibited from appearing on stage until the later scenes. Accordingly, by Brandom's lights, objective purport cannot be conceived until we have the main cast onstage at the beginning, to play out the first act. That cast of characters—assertion, inference, commitment, entitlement, and incompatibility—and the story they together perform—playing the game of giving and asking for reasons—are to be intelligible before the narrative of the later scenes, involving descendent characters. To break the analogy, perhaps, the actions of later characters are *not* to throw new light on those with which the audience are already familiar.

In the context of this structural design, Brandom cannot help himself to a practical phenomenon that seems essential to what assertion is: that we articulate (or try to articulate) how things are. I think we can profitably illuminate this point by reflecting on how assertions are actually individuated. We should not individuate assertions to be whatever have the grammatical form of a declarative sentence or a sentence in the indicative mood. As we saw in the last two chapters, not every apparent description is a description, even falsely, "of the facts". For instance, analytic propositions, in particular, are normative, and normative propositions are not intelligibly conceived as true or false of the facts; they prescribe and do not describe. "The chess king moves

---

[123] After "true" and "refers", those he focuses on are singular terms, anaphora, and *de re* ascriptions. See, respectively, Brandom (1994, chaps. 6, 7, 8, 9).

one square at a time", when used as a rule, does not (indeed, could not) describe a fact. But *real* assertions—i.e., sentences actually used to describe (in a number of different ways), in spite of grammatical form—are distinctive precisely because they are attempts to articulate how things are (or could be, or were, and so on). And that is, in practice, how we individuate them from other linguistic manoeuvres.

If, as I believe it to be, the foregoing thought is cogent, then Brandom is not even entitled to the notion of assertion in his Act One list of characters. For Brandom construes that notion to have no direct relation to the practice of describing how things are. This point is damning, because, as McDowell shows, it has two devastating consequences. First, Brandom himself is committed to the mutual dependence of inference and assertion: "Asserting cannot be understood apart from inferring. [And] . . . [i]nferring cannot be understood apart from asserting" (Brandom 1994, 158). Brandom's rationale for these claims is likely quite right. For example, to take each other as making genuine claims, we must presume each other to be able to show an understanding of those claims by telling us both their consequences and the circumstances which entitle us to them. But that is not all that the meaningfulness of assertions is bound up with; it is also the case that assertions describe how things are. So, granted that it is true that neither assertion nor inference can be understood independently of the other, when Brandom is not entitled to the notion of assertion, he is likewise not entitled to the notion of inference. Rather, as McDowell puts it, what was formerly called an assertion is simply a "move", and what was formerly presumed to be an inference is simply a "transition", in the deontic scorekeeping model (McDowell 2005, 128). Such is the risk of analysing one set of a vocabulary in terms of another: they may be inextricably related.

The other devastating consequence leads us nicely into the next subsection. If Brandom is not entitled to the concept of assertion, because the basic kind of move it is in the deontic scorekeeping model is not one which articulates how things are, then the model—the game of giving and asking for reasons—has no essential bearing to reality. This is not to say that people could not, really, play the game; it is only to say that prospective players would be making moves and transitions in a way that is self-contained. It is true that Brandom allows noninferential circumstances of application (perceptual in nature) and practical consequences (actions) to be legitimate transitions. But given how the notion of assertion, because representational, is actually inapplicable to the game, this does not change much: "nothing in the description of the deontic structure ensures that these pointers outside the game have anything to do with the sort of meaningfulness that would reveal the practice as linguistic" (McDowell 2005, 127).

Let us expand on McDowell's point with an example. Suppose a red ball crosses my line of sight, prompting me to speak. I respond with the move, "The ball is red". So I am entitled. But a

nearby scorekeeper wants to check I really understand it. So she asks me what else the ball must be, given that it is red. I respond, satisfactorily, with the claim "The ball is coloured". (For simplicity's sake, let us not consider also the exit from language conceived in terms of practical consequence.) According to the discursive scorekeeping model, this is a successful discursive interaction. It is the manifestation of conceptual capacities, which are exploited in inferential terms, even though the circumstance due to which I applied the concept was noninferential. McDowell is saying that this game completely fails to show that I actually was rationally entitled to the claim that the ball is red. Because, as Brandom sets this practice up, the circumstance of application has no *meaningful* relation to my utterance. After all, the move I make is not to be construed as one which is an articulation of how things are. So what ensures that I actually was entitled to this claim in the first place, and that it was indeed about the world around me?

§4.2 Observation reports

According to Brandom, even empirical conceptual content, the kind possessed by the concept "red", is articulated by the game of giving and asking for reasons. This is what we saw in the example above. For Brandom, although our responses to environmental stimuli—such as "That's red"—are in the paradigmatic case noninferentially prompted, he does not think that by being reliably so prompted they are thereby conceptually contentful claims. Instead, Brandom seeks to amend a reliabilist take on justification by specifying that observers reliably making reports are also participants in the game of giving and asking for reasons. Reliabilism in epistemology refers to the way in which the justification condition of traditional tripartite accounts of knowledge (where knowledge is equal to a justified true belief) is conceived, by affecting the way a true belief is formed: namely, by a reliable process. Brandom is unsatisfied with this, however, since what he calls "reliable differential responsive dispositions" do not suffice for knowledge (1994, 212). Rather, an utterance produced by such a disposition could only be a claim to know such-and-such if it were, first, a *claim* at all. And as we saw in the last subsection, an assertion is, for Brandom, inextricably bound up with inferential commitments. Accordingly, Brandom argues that an utterance such as "That's red", brought about by a reliable differential responsive disposition, is only contentful by dint of its possibly figuring here as premises and now as conclusions, in accordance with the material proprieties instituted and maintained by deontic scorekeeping

practices. Meeting both criteria, however, "That's red" is a bit of knowledge. Or so Brandom contends.[124]

As a model of perceptual knowledge, this fails to be intelligible. It is undermined by Brandom's reductivist commitment, and also therefore by what that entails about his notion of an assertion. Brandom introduces the inferentialist requirement on knowledge—that observers must be participants in a social deontic scorekeeping practice—in order to separate parrots and thermometers from humans as discursive beings (1994, 214). The problem here is that he has not managed to account for this discursive nature; McDowell's criticisms run deep. We could put the point like so: for Brandom, it is as if humans *qua* discursive are game-playing parrots. But that fails to take adequate account of the dimension of practice which includes responding to reasons given by the environment.[125] The subject who exclaims "That's red" is not recognisably making a rational move until we can understand that she encounters the fact that it (say, a ball) is red. Another way of saying that the subject must respond to reasons as such in her experience (as discussed in §2), is to say that the subject must encounter the fact, as the reason it is, in experience. In these terms, let us pose Brandom a question: Who is ever able to encounter such a fact in experience?

This is something of a trick question, given that Brandom does not have an account of experience according to which it represents how things are. After all, he does not deem the notion of experience as integral to the social deontic scorekeeping practice which he takes to exhaust the nature of discursive practice. (So at best the role of experience can be an afterthought.) Indeed, Brandom regards McDowell's emphasis on experiential content, in particular, as unmotivated, given that the rational constraint to our free thinking so sought after in *Mind and World* can, he argues, be played by (Brandom's version of) reliabilism.[126] Nevertheless, Brandom signals agreement with McDowell on an important point: that perceptual experiences must possess rational credentials (Brandom 1995, 248). And that is a requirement Brandom thinks he can meet:

---

[124] Brandom's account of observational reports is derived, he maintains, from Sellars (1956). See Brandom (1994, 214–17, 2002). This is another source of deep disagreement with McDowell, who reads this interpretation as "perverse" (McDowell 2010, 130). McDowell's own account is based partly on how he thinks Sellars ought to be read.

[125] I do think it is important to conceive rational responses to the environment as within the *practical* dimension; perhaps this is just another way of saying, in McDowell's terms, that conceptual capacities are operative in experience (after a certain training). Nevertheless, putting it this way coheres with the platitude that semantics must answer to pragmatics. The thought here is simply that Brandom's conception of practice is both too reductive and too inflexible to show how, in fact, semantics and pragmatics are interdependent.

[126] And Davidson's version too, according to Brandom (1995, 252).

taking it that the fact that the *reporter* is noninferentially disposed (in the right sort of circumstances) to acquire a belief with a given content provides a good *reason* for acquiring a belief with that content *oneself* (a reason the *attributor* of reliability could cite in *justifying* his own belief, even if the one whose belief it is could not). That reliability (in specified circumstances) as a reporter is likelihood of *truth* (in those circumstances) of reports ensures that the connection envisaged by reliabilists between reported facts and reports of them is *not* merely causal, but also rational. (Brandom 1995, 251–52)

We can bring Rose back to help explain this passage. Brandom's thinking here is that even though Rose, as a reporter, may not be aware that she is reliable when she says "That's red", a scorekeeper (an attributor) can be in a position to so tell. If the attributor can tell that Rose is reliable, it turns out, he thinks, that she does have a good reason (she is reliable at spotting red things). Further, her belief thus acquired ("That's red") has a content, which though it is noninferentially obtained, only counts as a belief in virtue of its ability to feature as a premise and a conclusion of inferences. The picture that builds up from this sketch, then, is one according to which Brandom allows perceptual experiences rationally to constrain beliefs, which are only beliefs in virtue of the inferentially articulated content they express. And, further, Brandom can then say that conceiving Rose's experience itself as that which provides the rational entitlement to her belief is unnecessary, since an external process, socially observed, already so entitles her.

However, there is a glaring hole in Brandom's account of observational knowledge. The reporter, Rose, still does not encounter the fact that an item is red, or the reason as such for her belief that an item is red, provided by her experience. It is, so to speak, completely foreign to her. The best she can do is come to realise that she is a reliable responder, but she need not do so to acquire her belief rationally (on Brandom's account). So the above-cited agreement between Brandom and McDowell is illusory: "My point about perceptual experiences is that they must *provide* rational credentials, not that they must have them" (McDowell 1998, 406). That is, it will not do simply to frame experiences as events which, in virtue of their reliably occurring, simply are warrants for belief. For one thing, conceiving rationality like this will not soothe the transcendental anxiety which first brought about the concern to find a way to reconcile receptivity and spontaneity, since the passive component of experience on Brandom's conception fails to show to a subject that things are thus and so. The fact cannot be in view as a fact for her.

Worse, for Brandom, is that the attributor in his story needs also to be an observer. The attributor, *qua* observer, must observe both that Rose judges that such-and-such, and indeed that such-and-such obtains. In light of this, it is completely unclear how the attributor can attribute, and no successive iterations of extra scorekeepers will resolve this issue:

Brandom thinks that by describing interactions among a multiplicity of individuals, who, considered individually, are not intelligibly responsive to anything as reasons for their responses, and hence do not intelligibly have anything in view, we can somehow make sense of their having in view, after all, one another's performances, and objective items that somehow come into view for them as subject matter for those performances and for counterpart performances of their own. I think this smacks of magic. (McDowell 1998, 408)

So we have to come back, eventually, to a subject's encountering the facts. But by doing so, we have to give up the thought that empirical conceptual content can be specified in purely inferential terms.

That is, we acknowledge ultimately that the practice of rationally responding to the environment is partly constitutive of the content of such empirical concepts as "red". It remains true, of course, that one should not take a subject to be making a genuine claim about her environment unless she can participate in reason-explanation. But there is no fundamental form such explanation must take. It is sufficient for a subject to justify taking in that the ball is red because the ball is red. It is a very good reason. The reason must not be supposed to be available independent of the ability to make relevant inferences with the concept of red, but this is just to insist on McDowell's original view in *Mind and World* that receptivity (passive experience) and spontaneity (active thought) work in concert.

§4.3 Conceptual content

We arrive at a place, then, where we must notice that empirical conceptual content is not purely representationally articulated, as might be suggested by a subscriber to the Given; not purely inferentially articulated as is suggested by Brandom; but, at least, both inferentially and representationally articulated. It seems to me that we should say that conceptual content is *rationally* articulated, because norms of rationality clearly govern the representational dimension of our discursive practice, which if the foregoing is correct cannot be displaced by purely inferential norms. If one's experience shows things to be thus and so, given a certain training, one ought to take it that things are thus and so. Brandom eschews such a view because of a commitment to an untenable reductivism—one which renders the practices of assertion and inference unrecognisable as discursive, and which maintains an implausible semantic thesis.

There are, presumably, indefinitely many norms which govern what we ought to say, given a discursive training. They are indefinite because there is not a settled number of ways to continue using concepts. Indeed, norms change over time, often in response to scrutiny. (And often without

our awareness, as we see very soon in the last chapter.) They are not only inferential, but they are all rational, since to use language is to deploy concepts which take on a distinctive logical shape through the way that they are used. Sometimes, normative propositions express the meaning of a term; a bachelor is an unmarried man. Other times, they will express proprieties of inference; if today is Tuesday, tomorrow is Wednesday. Still other times, they will express the rational entitlements which experience of the world provides us with; if the ball is red, "the ball is red" (is true). None of these tokens are descriptive, because they are expressions of normative, conceptual, analytic truths.

The conclusion I want to draw from the epistemological discussion, then, is that once we acknowledge that experience itself provides a subject with noninferential reasons for belief, then we can find a suitable way of reading this remark: Conceptual content is rationally articulated. This simply means that to explain what a concept is, we must detail its rational import. A good deal of that will involve the inferences into which it enters, but some of it must reflect what it concerns, in the world, if it does.

For example, the concept "red" is not illuminated for all purposes by listing the familiar inferential norms, "If something is red, then it isn't green", "Red is a colour", and so on. It is also illuminated by pointing to *that* colour, in a series of samples. Ostensive explanations such as these therefore play a grammatical role in our practices. (Perhaps this is unsurprising from a certain view; Wittgenstein (1953, §50) already regarded samples themselves as instruments of the language.) And that is precisely because such performances are contentful; the person who is learning colours is learning that that colour is to be called red. As before, we should always be wary of the reductivist impulse. In such circumstances, we should not abstract the practice of ostensive explanations above the practices with which it is related, and then construe it as antecedently intelligible of any such practices. Instead, when the philosophical worry calls for it, we must recognise the background knowledge one must have in order even to recognise pointing gestures, for example. Moreover, we must see that an inferential role for "red" is also but not only in view.

There is, then, only a superficial similarity between the views of Brandom on conceptual content and myself. The threat was that by construing grammatical propositions—analytic truths—as normative and thereby as partly concept-forming, I might have implied that inferential relations constitute conceptual content in a reductive fashion. This is what Brandom does by arguing that material proprieties of inference are unexplained explainers in his order of semantic explanation. We have seen what is wrong with that reductive way of elaborating conceptual content by overviewing McDowell's critique of Brandom: it leaves out the representational dimension of experience. We can, of course, restore that dimension simply enough, but only by resisting the

temptation to conceive of empirical content in some uninterpreted way—in the manner of content-bearing sense-data, say. Instead, we must, as it were, interpret content; we have to turn the Given into the given. That requires us to acknowledge how our mature experience of the world is necessarily conceptualised, in the sense that without a background normative training, expressible as analytic truth, we cannot experience the world as conceptually graspable, or encounter the facts.

Chapter Five
Analyticity and conceptual change

§1 Norms under question

The picture I have so far insisted on gives pride of place to normativity as it manifests in practice. That is, the norms that partly constitute meaning, or conceptual content, are often implicit practical principles which language users develop and adhere to over time. Their usage morphs into slightly new forms, of course, and sometimes more novel uses take hold and become normative, or nearly-normative. Norms are more explicitly brought into discussion when parties need to clarify their understanding. Thus, an ability to operate according to a norm is a good first place for estimating a speaker's conceptual grasp in a particular pocket of discourse, in which specific purposes are pertinent. Indeed, sometimes we require other speakers—especially in educational contexts—to demonstrate openly the degree and nuance of their understanding. As such, norms may perform a kind of epistemological role: one cites a norm to support one's using a word in a particular way. Significantly, and with respect mostly to the last chapter, the norms which we cite are not only inferential in nature; at times simply pointing to a sample is sufficient to show the requisite level of understanding, rather than making inferences.

As should be clear by this summary, I have found it necessary to relax any insistence that norms could "fully" articulate conceptual content. Might this be deemed a weakness in my account; could something else supply concepts with the content that expresses them? Well, another option—and it is certainly a popular one—is to bring in the causal dimension. One might suggest, in the manner of Kripke (1972), that concepts derive their content from links of usage causally related via long chains back to grounding baptisms. However, the discussion of the last chapter should make us immediately doubtful of the normative contribution that could be made by any causal process in this manner described.

I have said that norms perform a justificatory or epistemological role by being cited in support of one's concept-employment, including ostensive explanations and definitions. It is a broader but equally relevant point that only what is conceptual can perform this role. Recall, we saw that for a subject to be rationally entitled to her ostensive judgement "This wall is green" she must actually encounter *that the wall is green* in her experience (she must see the wall *as* green). A speaker seeming to make this judgement, without having had relevant training in using representational and colour concepts, is not in a position to so encounter the wall as green (by contrast, she can encounter the wall as an obstacle, in a sense). Its greenness is not something she is in a position to articulate, and

neither can this greenness, of its own, make her speech pattern "This wall is green" rationally responsive to her environment. For the only kind of thing which can provide rational credentials for a judgement will be, like judgements themselves, conceptual in structure; anything not so structured is brute, external to reason, like the stormy wind that topples over a tree. To believe otherwise is to subscribe to a form of the ever-mythical Given.

With a fear of the Given in mind, then, it is prima facie not possible for some causal chain of usage to articulate conceptual content—even for empirical concepts which seem to apply directly to the world, such as "green". For there is no mechanism of justificatory support available here; one cannot cite distant and questionable baptisms to justify one's present usage of a term: what makes *this* employment of "bachelor" the correct one, in this pocket of discourse, is (say) the norm that a bachelor is an uncommitted male.[127] None of this is to downplay causality: doubtless, the spread of a term is something that could be charted causally, if one had access to all the conversations made and writings ever produced. (Perhaps, in a looser sense of "cause", lexicography just as much traces normative usage as it does the causal links between those usages, in etymological and historical terms.[128])

But should we not still worry that my account of conceptual content thus far allows no role for causal processes, and, perhaps relatedly, no available way to chart the *progress* of concepts? More generally, what seems to be missing from the account so far is any idea of how "the world" can contribute to conceptual content (if not by some crude causal chains). I have indeed been at pains to point out that language is intentional—that we do as a matter of course represent the world—and so even though the world and its contents cannot directly give our words meaning, in the manner of the Given, it must play a significant role. A concern is emerging, then, that the pockets of discourse I have discussed nonetheless float free of the world—not in the way that Brandom's game does, but simply because the only norms I have welcomed into the fold which concern the world are ostensive, grammatical devices, within the "space of concepts". Even if, as I contend, conceptual content is not *given* by natural elements, do not such elements still contribute to how that content arranges itself in the first place?

Mark Wilson (2006) offers a set of views about language and the world which both conflicts and intersects with my own account as just sketched out. His magnum opus is sprawling, wide-ranging, and ambitious—it cannot be adequately summarised. But for my purposes, a key lesson

---

[127] Causal theories of meaning tend to have trouble with linguistic forms that are not names, in any case, such as predicates like "is a bachelor".

[128] Incidentally, the etymology for "bachelor" remains uncertain, though its usage history is fairly well documented (OED 2017).

Wilson hopes to impart, and which I need to scrutinise, is that our ability to understand concepts is far poorer than most philosophers estimate: "We need to frame, I think, a far more mitigated appraisal of our capacities to anticipate our linguistic futures" (Wilson 2006, 11). Wilson is no Wittgensteinian, though one may recognise certain affinities (see Wilson's (2006, xx, 279, 566) own account of similarities and differences.) The mantra "Back to the rough ground!" (Wittgenstein 1953, §107; cf. Wilson 2006, 17) surely applies, as an injunction to focus on the use of terms—though, in Wilson's case, the rough ground tends to be applied mathematics and the physics of macroscopic phenomena. Indeed, an emphasis on normative use as such is missing; rather, Wilson focuses on descriptive practice and the strategies employed by various technicians to overcome practical problems (2006, 10), often when existent norms break down.

Accordingly, two of Wilson's claims need assessment: that our grasp of concepts is much worse than we think it is; and that the norms governing usage do not ultimately have much of a constitutive role, unlike how I have advertised, especially for predicates which "wander" away from original domains of usage. For an example of this second claim, Wilson examines the concept of hardness, and finds that what pushes usage one way rather than another has little to do with normativity:

> Whatever "rules" our linguistic peers might have originally taught us in regard to "hardness" prove quite beside the point, soon overcome by the impertinent particularities of metal, ceramic and plastic. (Wilson 2006, 344)

The fourth section of this chapter takes up the challenge of responding to this second claim.

To get a handle on the strategic development of some predicate usage, Wilson employs a largely novel vocabulary and imagery, only some of which I will relay here. Typically, the usage of a classifying predicate such as "is hard" is mapped onto overlapping "patches"; the predicate may sit, in his metaphor, atop these patches, which themselves straddle above the world of properties or attributes. Wilson points out that these many patches of "is hard" usage find themselves arranged in a messy manner, overlapping at their boundaries in some cases and not others, and in fact likely not attaching to the same properties in each employment—without, for all that, being especially vague or difficult to apply. Wilson describes this structural image of descriptive practice as an "atlas" or "façade",[129] arranged in "patchwork" style. The story Wilson tells of how these patchwork structures accrue new patches and are navigated between is rather compelling. Further, it is not too dissimilar to what I have gestured at, in much less detail, with my "pockets of

---

[129] I will be dropping this latter term, for simplicity.

discourse" talk. But, as foregrounded, a prima facie difference is that Wilson finds a much less prominent role for normativity in his narrative (Wilson 2006, 343–44).

Why does Wilson downplay the importance of normativity in discussing conceptual content, employment, and behaviour? If Wilson is right to so reduce its significance, then my own picture—somewhat similar in its contours—is in dire need of reform. Further, I shall have to characterise philosophy itself in a quite different manner if conceptual content is not usefully articulated by norms but by something else. Fortunately, I believe Wilson's reasons for limiting the contribution of norms to conceptual content derive from a fear of too rigid normativity, from which my own picture does not suffer. (There is a structural similarity in this respect, although their projects differ greatly, between Wilson and Williamson.) Nonetheless, Wilson's framework for examining concepts gleans insights that I wish to elaborate on for the purposes of this thesis. And so overviewing this work presently will be instructive.

§2 Semantic finality and the patchwork structure

Wilson makes use of so many examples—from disciplines as wide-ranging as applied mathematics, mechanics, geography, atlas-making, various industries, folklore, musicology, applied physics, and biology—it is hard to settle on just two or three to demonstrate the rough picture. We can best bring about what is important by first outlining the commitments of those to whom he objects.

Wilson groups together a number of theses under the heading "the classical view [or picture] of concepts" (2006, xiii-xiv; 7).[130] He takes Russell to be the most "perfect representative" of this view (2006, 5), though insists it finds expression in a list of thinkers as diverse as Frege, Carnap, Kripke, Nelson Goodman, neo-Kantians, and communitarian kinds of Wittgensteinian, sometimes in "neo-classical" form (e.g., see Wilson (2006, 83–84)). Further, the body of views has changed in some ways over time, so that not all its theses would be accepted by a good Russellian, though the changes seem like natural extensions (2006, 5). Wilson finds this group of views to stem from everyday thinking, unhappily titled "ur-philosophy"—vaguely philosophical perspectives which do not encounter much pressure in normal environments, but which still dictate patterns of thought in headier or more practical matters, where they place in us an undue confidence in our conceptual abilities (2006, 17).

One way in which this overconfidence manifests in philosophical thinking, according to Wilson, is in the views of the "ordinary language movement", which maintains that we

---

[130] Although he does not attribute to any one thinker all the theses, or insist on their exact form, Wilson does provide a list of 44 interconnected theses at the end of his third chapter (2006, 139–46).

learn complex, implicit rules from our linguistic tutors that restrict "concept" and "attribute" to finer circuits of proper application. If we would only attend to these rules, it is argued, we should be able to prevent language "from going on holiday"[131] in the manner that leads to errant philosophizing. . . . [S]uch projects rest upon an untenable view of language insofar as they demand a foundation in the notion that "our linguistic training tells us how to use notions like 'concept' properly." (Wilson 2006, 18–19)

Notably, Wilson does single out J. L. Austin as a philosopher with whom he shares methodological objectives and sympathies. However, it strikes me that perhaps more than any of those sometimes portrayed as "ordinary language" philosophers (which is not, to my mind, a healthy portrayal), Gilbert Ryle is the one whose methodological tenets most resemble some of Wilson's key approaches. This is something I shall detail further in §4 below. At present, we should note that Wilson is here denying that mastery of a concept is something determined by learning rules early on in learning a language. Rather, his view is that concepts have an open-ended and surprising range of application, which often we discover not by reflecting on the norms governing early concept-usage, but rather after implicitly and explicitly adapting our usage to a stubborn, unforeseen world. Moreover, it is not just concepts of everyday phenomena which come under Wilson's scrutiny, but also terms of "conceptual evaluation" themselves: "concept", "grasp", "notion", "property", and so on.[132] Thus Wilson is arguing that linguistic training fails to instil within us failsafe rules for applying evaluative terms too—for knowing what it often amounts to in saying, e.g., "Tommy doesn't quite *grasp the concept* of motion". On Wilson's account, our inadequate training thereby skews not only our ability to determine the "right" way to employ a given concept, but also the way to get a measure of correctness concerns themselves.

This idea that we attain a full grasp of conceptual contents generally when young, as part of learning the language, Wilson dubs "semantic finality" (2006, 19). As a core element of the classical picture of concepts, he ascribes it not only to ordinary language philosophers, but to other thinkers who help themselves to this picture. What leads us to assume semantic finality, he claims, is a desire to underwrite basic features of language such as communication and understanding; we hope fixed contents can ensure that we relay what we intend, and that our audience can interpret us correctly. He finds another motivation in the felt need to explain the creativity of language, our ability to state and understand limitless new propositions, in terms of a finite compositional structure (Wilson cites Ray Jackendoff convincingly to this effect—2006, 19–20).

---

[131] Wilson cites Wittgenstein (1953, §38) for this remark.

[132] See (2006, 2) for "conceptual evaluation" talk.

Although broad strokes, the thesis of semantic finality does strike me as accurately describing semantic views shared throughout different periods. One can certainly recognise, for example, the motivation to find units that cross a chasm between speakers, allowing us to understand each other, in the writings of Locke (Locke 1689, bk.III, ii, §1).[133] More exactly, one finds expression of something very similar in Russell's descriptivist treatment of proper names (B. Russell 1956, 243), and a common way of reading the senses attached to proper names in Frege (1892). Moreover, according to semantic finality as set out by Wilson, while conceptual content may be slightly tinkered with, ultimately any new learnings are not absorbed into the content proper, and rather remain external to core conceptual content. This certainly resembles how the analytic-synthetic divide is often described, where core content ("All bachelors are unmarried") is analytic and externally learnt matters ("Some bachelors are miserable") are synthetic. As we saw in the first chapter, the notion of analyticity is itself pertinent to many branches of historical and some contemporary philosophy; as such, semantic finality may be a fair generalisation of a theme that has run throughout analytic philosophy in recent times at least.

Indeed, that semantic finality and common interpretations of analyticity are similar should prompt us to anticipate how the picture of normativity defended in my thesis will not fall prey to concerns about semantic finality. What is so wrong about semantic finality, anyway, according to Wilson? Simply put, he contends it is an inaccurate account of conceptual content:

> As we continue to work with our words past our hypothetical date of finalized capacity, virtually every term of macroscopic evaluation becomes subject to subsequent shaping pressures for which our training has left us unprepared. In compensation, subtle correctives and barriers creep into our language, often quite unnoticed, with the net effect of turning our classificatory concepts in quite different directions than we originally pictured. (Wilson 2006, 20–21)

That is, according to Wilson, concepts used to classify items of experience—say, in respect of their hardness or colour—are "shaped" in the course of their usage by pressures that we may not even recognise. And so, to insist that conceptual content is, so to speak, settled in advance is to sell short the capacity of our terms to evolve. Worse, he thinks, is how obstructive to learning more about the world this view can be. For he claims that some significant advances in, *inter alia*, applied mathematics, gear-wheel mechanics, and classical physics would not be possible were thinkers to cleave too strongly to the content of concepts as they had been taught them.

---

[133] On the classical picture's founding in Locke and Hobbes, see Wilson (2006, 101).

An example will serve to make clear exactly why semantic finality, according to Wilson, cannot hold for all our concepts. Fairy tales often feature rainbows which display the qualities of some physical arch, as though they were scalable and locatable. As Wilson notes, we might even consider it criterial for the notion "rainbow" that a young child be able to understand arch-like employments of "rainbow" (2006, 21). Growing up, however, we find this picture false and the stuff of fantasy; instead, we learn that the rainbow is a product of raindrops irradiated in a certain manner, and usage which portrays rainbows as continuing to exhibit the form of material arches we will classify as errant.

There are two points to make about this example. The first is that whatever linguistic competence was instilled within us when young with respect to the concept of a rainbow, plainly it did not inform how the predicate's application should change, perhaps like a seedling to a tree. The predicate's changes rather become necessary when we "accommodate" to "real world contours" (2006, 23). The basic message here is that semantic finality does not hold for the apparently pedestrian concept, "rainbow". And Wilson urges us to imagine the consequences of misconstruing the contents of other, more important concepts, such as "force" and "hardness", because of a foolhardy commitment to semantic finality: "our buildings fall down and our knife blades dull at inopportune moments if we augur their conceptual contents wrongly" (2006, 22).

The second point, though later picked up somewhat via the metaphor of "latent DNA" (2006, 459), is something which Wilson could emphasise more (I discuss this thought in greater detail in §3.2). Namely, we do not scrap the sorts of test we put before the child; within certain circumstances the concept remains *properly* employed in "arch"-like manner.[134] On Wilson's part, he highlights here the purpose of concept-evaluative vocabulary such as "concept" and "attribute": by means of them we can "monitor and correct our usage", and so our notion of conceptual grasp itself "must display considerable sensitivity to the maturational level of the speakers we attempt to evaluate" (Wilson 2006, 23). That is, we need not conclude that a 7-year-old misunderstands the concept of a rainbow just because she responds to a rainbow's apparent arch-like features.

What Wilson hopes to come through in this example is how our normative classificatory practices ultimately play second fiddle to Nature with a big "N". Indeed, Wilson exhibits a trend to describe predicates and concepts by use of metaphors depicting action without human

---

[134] It is unclear from later passages just how much Wilson would be committed to the view as stated here. Nonetheless, I think it is an important one. Wilson's metaphor is intended to capture how old "directivities"—normative standards, more or less—can help new usages nucleate. But it strikes me that Wilson could, if so inclined, make this point without recognising the enduring propriety of those old directivities, much as one could recognise the influence of some Roman law on contemporary common law, without believing Roman law to govern some present jurisdictions.

intention: they get "pulled" onto different patches of usage (2006, 257), and thus "drag" properties (2006, 175); they are "shaped" by external pressures (2006, 135); they evolve (2006, 8), have careers (2006, 23), and so forth. The picture Wilson slowly paints is one in which predicate usage conforms to the obstacles which "Nature" stubbornly presents; we conform to its ways, often adeptly and unawares. Sometimes he does speak of our cooperation with Nature (2006, 235–36), but this is not by means of conforming to the rules which we think govern the use of the concept. Indeed, he argues we simply cannot be confident that some stock of conceptual content remains firm and directs or prescribes future usage, for there always remain "untested domains of application" (2006, 22).

Take "rainbow" again. Even knowing that a rainbow is formed from irradiated raindrops, one might have thought it impossible or nonsensical for a rainbow to "lie on its side" (2006, 22). Yet, as we find out, this becomes a well-known phenomenon in certain conditions (2006, 434–35), which thus cannot be ruled out in advance by some invariant conceptual content. Wilson offers a more striking example of untested application domains: the colour of rubies on Pluto. Being cold and poorly lit, it is plausible, he contends, that conditions on Pluto alter the crystal array of rubies. And since rubies' hues "depend sensitively on scattered color impurities in their matrix", Wilson thinks it possible that a Plutonian ruby could possess a quite green hue—as revealed in a time exposure photograph (so that humans could view it despite Pluto's low light levels) (2006, 231–32).

Were this "green ruby" scenario to obtain in conditions such as Pluto's, it is unknown how we would adapt. Wilson points to existent prima facie inconsistencies: that we allow a heated corundum[135] to classify no longer as a ruby by virtue of a wrong temperature change, yet we also allow that a ruby is red even at night with the lights off. Seemingly, then, the "directivities" (2006, 15) or guiding normative standards that we follow change as we skip over into different patches. (Wilson's (2006, 95) term "directivity" is used "as a non-technical means for capturing the loose bundle of considerations that we might reasonably cite, at various moments in a predicate's career, in deciding how the term should be *rightly applied*.") Accordingly, from Wilson's perspective, we cannot even ensure that our linguistic training for "red" or "ruby" equips us with a core conceptual content, able to classify any item it meets. Rather, "we will have merely assembled preparation adequate only to a narrow, local slice of the universe" (2006, 232). To assume otherwise is to betray "tropospheric complacence" (2006, 55): to believe that conditions of application remain relevantly

---

[135] Wilson mistakenly refers to the minerals of sapphires and rubies as *beryls*, though this is incorrect: beryls compose, *inter alia*, emeralds and red beryls, depending on their array of impurities, but not rubies and sapphires, which are varieties of the *corundum* mineral.

similar come what may, such that we will always have an easy understanding of what something's being red or icy will amount to in any environment. (Incidentally, Charles Travis (2008a, 14) says that criticism of such a view is "a basic point of Wittgenstein's methodology".)

So much is what Wilson hopes to have shown from the above thus far: that the classical picture of concepts, sprung as it is from everyday thinking, inclines us toward some complacency; we assume that the circumstances in which we employ concepts remain the same as we go on, and we assume that the normative standards to which we adhere in so employing concepts remain perfectly adequate. Yet he points out that both of these assumptions turn out to be ill-founded when we look at actual cases, and so the classical view, at least in respect of its semantic finality thesis, needs to be overturned. However, he does not recommend inverting it totally; indeed, he thinks postmodernism emerged as a too extreme repudiation of certain classical themes (2006, 3), and thinks softer though still misguided rebellions are at the heart of Quine's holism and aspects of Wittgenstein[136] (for instance, the emphasis on rules and practices as they pertain to meaning). In that respect, then, Wilson regards classical thinking about concepts as on the right track, albeit too stiffly:

> Forced to choose between exaggerated mistrust and blind acceptance of every passing claim of conceptual authority (even those issuing from transparent charlatans), we should plainly select gullibility as the wiser course, for the naïve explorer who trusts her somewhat inadequate map generally fares better than the doubter who accepts nothing. (Wilson 2006, 3)

I sympathise with Wilson's outlook here. He is unwilling to withhold trust, completely, in conceptual "authority"—i.e., authority over the proper, normative application of concepts. And this resembles my own response to Williamson. Recall §4.2 of Chapter 3, in which I held that one cannot say, on a whim, as per Wittgenstein's example, "bububu" and simply mean "If it's not raining outside, I shall go for a walk", because language is not totally free; although its norms spread diffusely, grow branches, and change, it remains a normative system. (One might say that the freedom characteristic of language is a freedom conditioned by limitations, in this case normative ones; in which case, "total freedom" of language is no freedom at all.) Similarly, in Wilson's narrative, by submitting wholly to mistrust when it comes to matters of conceptual authority, we would be depriving language of something essential to it: we "would render daily life patently unworkable" (2006, 3). Indeed, a central point I make against Williamson in Chapter 3 is,

---

[136] It should be noted that Wilson still regards Quine and Wittgenstein as both central figures in seeing the faults of classical thinking in the first place (2006, 10).

in Wilson's terms, that he provides no tools with which speakers could evaluate for misunderstanding of concepts, and yet we plainly do have such tools and implement them well enough.

Set against this backdrop, then, it is difficult to locate precisely the point of disagreement between Wilson's picture and my own. He wants both to commend our ability to resolve matters of conceptual misunderstanding and to resist overstating that competence. So much I can agree with. Yet a consequence he wants to extract here, among others, is that governing linguistic rules and such do not primarily contribute to conceptual content. Moreover, on his view, we do well to ignore these rules to bring about successful descriptive practice. That is, on Wilson's view, rules we learn for the application of concepts in key cases must be overturned; we must forge ahead, norms be damned, in order to describe correctly the objective world.

I think Wilson is simply wrong on these latter points: language is, on my picture, irreducibly normative, and the methods used for conceptual evaluation in our linguistic practices by and large indicate that norms do underwrite conceptual content. The challenge, then, is to reinstall normativity's seat at the linguistic table, yet to do so in a way which respects the ability of technicians to change the way they apply key concepts in their investigations, in response to worldly prompts. Wilson's reaction to one of his favoured examples is especially useful for taking on this challenge.

§3 Weighty concerns

Consider predicates such as "weighs x pounds", "is weightless", and so on. On Earth, students learn to distinguish the notion of weight from mass, since they pick out different properties: mass is a measurement of matter comprising an object, whereas weight is the impression that mass makes as it is pulled by a gravitational force.[137] Accordingly, students learn that to apply "weighs x pounds" to a person or object, the measurement is relative to the exerted gravitational force upon that person's or object's mass: what Alfred weighs on Earth (180 pounds, say) is not what he weighs on the moon (30 pounds, say), given the moon's much weaker gravitational field. Moreover, Alfred's lightness on the moon is reflected in the other salient directivities or norms which speakers implicitly follow. For example, we learn that if something is heavy it is harder to lift (or carry) than something lighter; that it will hurt more if dropped on one's toe; and so on.

---

[137] This is how Wilson presents the discussion (2006, 328), although it seems to be incomplete. According to Einstein, mass is the aggregation of energy (potential, kinetic, thermal, etc.), not matter. See Markus Pössel (2010) for more on Einstein's view of mass and energy, and Einstein (1916) for an accessible overview of the following topics, especially those in §3.1.

Likewise, when Alfred is on the moon, he is easier to lift than on the Earth, and he would register 30lbs on a suitable scale. Finally, when we see astronauts in space stations, slowly floating around, we judge them to be weightless; they can easily carry heavy machinery which on Earth weighs several times their earthly weight, and they would register 0lbs on a scale.

Accordingly, we have norms by which to ascertain something's being heavy or light and what follows from such judgements; further, we learn the norm that to apply the concept of weight properly, we must distinguish it from an object's mass and calculate weight according to the gravitational field in which the object resides. So far, seemingly so consistent. The problem Wilson poses crops up, he says, when we think about Alfred's different weight assignments collectively and sequentially over points in a larger expanse of space. We can tabulate Alfred's weights onto a single chart, jotting down his weight at separate locations: as he rockets out of Earth, past its atmosphere, through space, towards the moon, and finally lands on the moon. Say he weighs 30lbs on the moon; working backwards, at some point, just where the gravitational pull of the moon and the Earth cancel each other out (one of the so-called Lagrangian points), the exerted force upon his mass will mean his weight is 0lbs. But as he continues to travel back towards the Earth, he starts to pick up in weight, relative to the increased force exerted upon him. At some point, about 250 miles from the Earth's surface, Alfred will bump into his space station colleagues. There, we have previously been willing to accept he and others are weightless, but in so doing, we have not stayed firm to an earlier principle: that an object's weight is a function of the gravitational force exerted on its mass.

For, in the orbiting space station, the gravitational force exerted on Alfred is 88.8 percent of what it is on the Earth (NASA 2015). Accordingly, he should weigh around 160lbs. Here, we encounter something Wilson calls "multi-valuedness" (2006, 329). Inspired by the way applied mathematicians describe the behaviour of analytic functions (2006, 316–17), Wilson applies the term to a wide range of non-mathematical predicates.[138] Crudely, the idea is that, in following distinct standards for applying some predicate, we end up jointly endorsing—and seemingly without error—competing, contrary values or claims: in our case, that Alfred weighs both 0lbs and 160lbs when at the orbiting space station. Wilson's explanation for this multi-valued phenomenon is that we have tacitly allowed ourselves to be directed by different salient features of the scenarios we describe, such that we actually end up aligning the predicate "weighs x pounds" with different properties:

---

[138] The mathematical usage applies to this kind of case: if not sensitive to context of application, calculations with complex numbers may seemingly prove "+2=-2" (Wilson 2006, 312).

In the face of so much apparent classificatory reward and so little discouragement, we remain well satisfied with our "weighs zero pounds" judgments. Thus the powerful inferential and classificatory directivities of our folk physics reasoning standards silently shape our space station employment of "weight" in a parochially adjusted manner, despite the fact that the predicate is thereby *dragged away* [. . .] from its erstwhile attachment to [the property] *impressed gravitational force*. (Wilson 2006, 330)[139]

That is, the normative standards we incline to stay true to, in coming to the "zero pounds" judgement, are those which we apply in terrestrial contexts, or "patches" covering a terrestrial domain, in Wilson's terms. And in jumping between most patches for "weight", Wilson suggests, we have no problem remaining consistent with the principle for deriving an object's weight, which connects "weight" to the property *impressed gravitational force*. Yet on one kind of jump between patches of usage, we are somehow tricked into applying "weighs x pounds" to a different property: "*work required to move* x *relative to local frame*" (2006, 330).

Now, Wilson's account of what happens here is not as crude as it might seem. He is not saying, for example, that the weight predicate is simply applied wrongly whenever we describe astronauts on the station as weightless. Rather, he argues that we do have rules for extending the application from the territorial patch of usage covering the Earth to the patch covering the station. By Wilson's lights, such "prolongation rules" are successful and do not create problems considered on their own. Likewise, we have equally useful rules for prolonging the predicate's usage across the domain of space stretching from the Earth to the moon. According to Wilson, however, we simply lack a prolongation rule connecting the patch covering the whole stretch of Earth-moon space with the patch covering the orbiting space station (2006, 330–31).

Thus, Wilson represents the structure of usage as a series of patches connected by boundary rules to help us switch between patches. Sometimes these patchwork structures or atlases just exhibit, he thinks, multi-valuedness, and it is folly to try and "flatten" them out into more uniform usage—Wilson maintains that the best we can do, and indeed tend to do, is pay attention to the boundaries of these patches and work out where there are rules to prolong the usage, and where such moves are blocked. The weight example, however, is one for which he contends there is a plausible technique to so "flatten" out these different usages—and it is this I find most contentious: "we simply need to stick 'apparent' in front of 'weightless' every time we employ the phrase within a space station context" (2006, 382). To my mind, this is not as effective a trick as it would need to be to create some uniform application for the predicate; examining just why

---

[139] Note, Wilson uses his own stylistic scheme, in which a property is depicted by a phrase in boldface italics. I have chosen not to reproduce that scheme here.

demonstrates that Wilson stops short of the more radical consequences of his thesis, which he ought nonetheless to endorse. Moreover, as I show, there is good reason to do so.

Preceding my argument to show that we cannot reasonably hope to "flatten" the predicate usage of "weight", it is noteworthy that there is a technical explanation why one should sometimes refrain from calling space shuttle astronauts "weightless"—at least given a Newtonian picture of physics, and specifically gravity. The International Space Station (ISS) is actually travelling at approximately 17,500mph in its orbit around the Earth (NASA 2015). We can describe the ISS as maintaining this orbital speed in order to balance the gravitational pull against the centrifugal force thus generated by its horizontal motion. In doing so, the station is, so to speak, "falling around the Earth"—it is in gravitational free fall, much as one would be if airborne in a fast elevator moving downward from a great height, or (allowing for air resistance) on certain similar fairground rides. When in free fall, one has the sensation of weightlessness—of what would be experienced at a Lagrangian point between the Earth's and the Moon's respective gravitational fields, say, where Alfred's weight is zero pounds. We can see, then, why it is tempting (on a Newtonian picture) to say that Alfred is only *seemingly* weightless when orbiting the Earth. But this fails to do justice to two other normative considerations: our application of "weight" given, first, Einstein's general theory of relativity, and, second, other everyday normative standards which contribute to the content of "weightless" (and, indeed, to that of "seemingly" or "apparent").

§3.1 Einstein

Let us consider the first point. Wilson's discussion of the weight scenario clearly takes Newtonian cues; he does not spell out how "weight" might change in application once we take Einstein's work

Fig.1
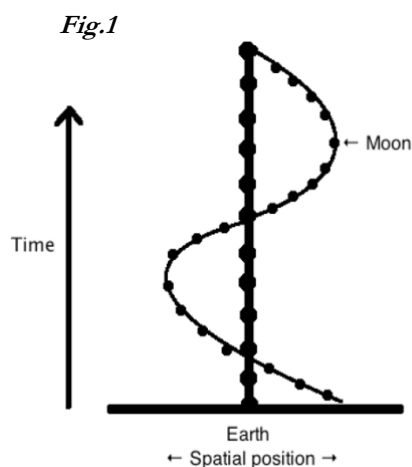
on relativity fully into account. Einstein's reading of the example, in lay terms, may be along the following lines.

We must reorient our conception of gravity, first and foremost. To see how, consider Einstein's notion of *spacetime*—he imagined that space and time, conjointly the "fabric" of the universe, are not dimensions which we can dissect apart. Having made this move, he then showed that what we incline to see as the gravitational pull of bodies could actually be the straight paths bodies take along the "curvature" of spacetime. Finally, given the curvature of spacetime and the paths of gravitating bodies along it, he was able to propose that gravity is not really a force at all, but simply a feature of spacetime geometry. Let us unpack these details.

It will help to look, by comparison, at a Newtonian way of representing gravitational orbits. Look at Fig. 1; it shows the orbit of the Moon around the Earth on a 2D spacetime diagram: keeping the Earth's position fixed, as time passes, the Moon's position (simplified to one spatial dimension for ease of presentation), curls around that of the Earth. According to Newton, if the Earth were not present, then the Moon would carry on in an inertial (i.e., uniform) path as in Fig. 2, or else be at rest. Newton thus conceives a force—gravity—as what makes the Moon orbit around the Earth when in its proximity, thereby pulling it back from its inertial path.

Fig. 2

Time

← Moon
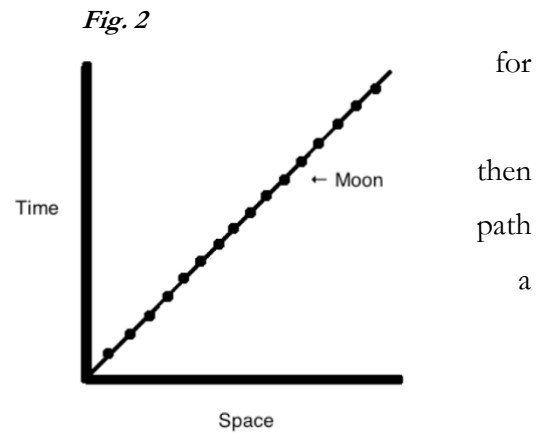
Space

But while this is an accurate model for most orbits, ultimately it proved weaker in predictive power than the model Einstein proposed. Einstein managed to confirm, instead, that spacetime is *itself* curved. This means that the orbiting path in Fig. 1 is already inertial, travelling in a uniform, straight manner; it somewhat resembles the line drawn in Fig. 2, then, because it takes this (nearly) straight path. The straightest possible line, or shortest possible path, is what is known as a "geodesic", which is analogous to an exactly straight line in 2D plane geometry, i.e., the shortest distance between two points. But as applied to 3D geometry—e.g., as on the surface of a sphere— the straightest line between two points will fall along a curve (just as it is on Earth); it will be the "shortest curve" between two points. However, in 4D or spacetime geometry, the shortest curve between two points is one that curves along *both* space and time. Unsurprisingly, we cannot easily conceive what this looks like except by relying, somewhat illegitimately, on the spherical model and bracketing the curving of time. But suffice it to say, most significantly, that any free-falling particles or bodies—i.e., those acted on by no forces except for gravity (which is not quite a force anyhow)—are properly described as taking spacetime geodesics (the shortest curves along spacetime).

Taking this into account, an early worry presents itself for Wilson's attempt to flatten out the "weight" predicate onto a patch covering just the property *impressed gravitational force*. For according to Einstein, gravity is not a force to begin with! In my example, the Earth's mass is intimately co-operating with the local fabric of the universe—i.e., reciprocally and simultaneously, the Earth's mass curves the surrounding spacetime and the spacetime "moulds" the mass. And so Alfred on the ISS, free-falling around the Earth, is travelling as straight as possible along the spacetime curved by the Earth's mass. This means that Alfred is not, as per Newton, only following the spherical curvature of the Earth by horizontally travelling fast enough to offset its gravitational

pull. Rather, he is simply following the spacetime geodesic which matches the curvature around the Earth. Therefore, Alfred is not forced away from the straight path he would have travelled in the absence of the Earth; he is exactly carrying on in the straightest possible line given the Earth's gravitational distortions.

One way to express what happens here is to say that although Alfred is accelerating around the Earth, he may as well be at rest in outer space. This is a rough statement of Einstein's equivalence principle, which has a direct consequence for how we think of weight. First, we have to understand so-called frames of reference. By Einstein's theory of special relativity, an astronaut in a space shuttle, with measuring instruments and a clock, freely floating in space (far from any significant bodies), is in an "inertial" frame of reference.[140] To qualify as being in such a frame of reference, the observer—an inertial observer—must be moving at a constant velocity (speed in a direction) or be completely at rest. Two such observers may see each other go past their spaceships; without anything else in the frame, both will see the other as moving past herself. But the moment one observer accelerates, she stops being in an inertial frame of reference. Instead, she becomes an accelerated or non-inertial observer. When he offered his special theory, Einstein held these two kinds of frame distinct. However, when he tried to adapt special relativity to account for gravitational forces, he was forced to conclude something strange. Namely, he postulated that an accelerated frame of reference is in effect no different than an inertial frame of reference so long as the acceleration matches the apparent force of gravity on the inertial frame.

Let us get a picture of this point via an example, by means of which we can see the relevance to the notion of weight. The pull of gravity is measured in Gs; the Earth's pull of bodies towards its surface is stipulated to count as 1G ($\sim$9.8m/s^2). If an astronaut in a windowless box is accelerating through open space at a rate of 1G, in any direction, her feet will be planted to the floor of the box, and she will experience the same feeling of "weight" in her limbs as she does on the surface of the Earth. Likewise, objects will fall to the ground at their usual speed when thrown (*mutatis mutandis* for a lack of terrestrial air resistance). Thus, Einstein asserted an equivalence between the rate of acceleration and the "force" of gravity, which may as well be regarded as acceleration. This is Einstein's famous equivalence principle, in layman's, non-mathematical terms. (Interestingly, as DiSalle notes (2016, §2.6), the principle was in part anticipated by Newton himself.)

There is here a slight complication with this equivalence: bodies do not take parallel trajectories (relative to each other) falling through space toward a body's mass (neither do they on the surface

---

[140] On frames of reference, see Pössel (2005).

of the Earth, but the effect is so minute we could not detect it). Rather, the bodies are falling in a *non-uniform* manner; further, the bottom of a body (relative to the Earth) is also falling at a very slightly faster rate than the relative top of the body. Such differences are accounted for under the description "tidal forces". So the equivalence principle implies that an accelerating frame of reference in outer space, with an acceleration rate of 1G, is qualitatively the same as the inertial frame of reference on the surface of the Earth, excluding tidal effects.

What does this all mean for weight and weightlessness? Recall that on the Newtonian conception of gravity, Alfred on the ISS is still under the force of 88.8% of the Earth's gravity, and so in spite of his seeming weightlessness, he actually weighs 88.8% of his weight on Earth (160lbs). That is because, for Newton, the horizontal motion of the aircraft is fighting against the Earth's strong gravitational pull. But on the Einsteinian conception as articulated in his general theory of relativity, there is no such gravitational force. Rather, the Earth, which in Einstein's theory is bending the fabric of the universe, is understood to instruct, as it were, the orbiting body to take a uniform line along the curvature of spacetime. Given the equivalence principle, Alfred's accelerating frame of reference (his perpetual free fall around the Earth) is qualitatively the same as any *non*-accelerating, inertial frame of reference in outer space, where even Newton would concede he is weightless. (The only remaining question here is whether to count the small tidal forces as in effect making Alfred have a little weight, in this scenario.)

So, if Einstein is right (and decades of physics agree with him), Wilson's worry about "multi-valued" applications of "weight" may be unnecessary. Perhaps a Newtonian picture tethers "weight" to *impressed gravitational force*, but Einstein argues against seeing gravity as a force at all to begin with. Rather, "weight" will tether to a property such as *acceleration relative to frame*, and it will always be such, whether at a Lagrangian point, the moon, the Earth, or the ISS. In my view, in fact, given Einstein's discovery, we do well to think of Alfred in outer space, accelerating inside a shuttle at 1G (in any consistent direction), as having weight. Only in the context of Newtonian classical mechanics can refusing to describe Alfred as being weighty be worthwhile. Indeed, there are ordinary concerns about what it is to have weight which we should adduce in favour of saying that Alfred has weight here, and that is the point in the sequel.

§3.2 Everyday directivities

What the Einstein discussion shows, I think, is the importance of what we earlier remarked upon: "latent DNA". Before, during, and after the Newtonian conception of gravity and physics more generally, persons needed ways to refer to objects as taking more effort to lift, move, and carry. There are doubtless indefinite reasons why one would require a concept of weight in everyday

dealings. Insofar as the concept becomes distinct from others, of course it acquires its own inferential directivities: criteria on the basis of which one can and cannot make certain inferences. It so happens that in the case of "weight", those directivities pulled us in a different direction than Newtonian physics deems proper. Or at least, Newtonian physics taken strictly. What Wilson usefully proposes is a way to understand the concept of weight as taking on groups of divergent directivities. According to Wilson, then, it matters on which patch of usage we stand when we pronounce whether Alfred is weightless or not.

That is well and good, but let us not forget that at other times Wilson is keen to show that the undulating structures of usage he terms "patchwork" can sometimes be "flattened out". In the case of weight, he contends that we can do so by appending "apparently" to "weightless" to account for those patches involving gravitational free fall as in orbit. In this subsection, I offer more reasons why this solution is insufficient. In brief, it fails to account for the inferences legitimately warranted by the everyday directivities.

Recall that one learns to make certain inferences to and from a judgement that something is heavy (to speak, temporarily, in a Brandomian tone). In a terrestrial context, e.g., if the rocket attachment weighs a tonne, Alfred cannot lift it. What happens to these implicit norms when we move to new, extra-terrestrial domains? The issue is complicated. If, outside the space shuttle, Alfred refuses to lift the rocket attachment and pass to his colleague on the grounds that it "weighs a tonne" (or 88.8% of a tonne), we should have good reason to presume he is either joking, sulking, or has misunderstood.[141] That is, in this domain of application, the impressed gravitational force (given a Newtonian picture) on the rocket attachment is not a salient feature of its being weightless: rather, the purposes of this task require that astronauts pay attention to other aspects of something's being weightless. In this case, that the object can be lifted and carried with ease.

Furthermore, the effect of weightlessness produces notable changes to the human body:

> During evolution on Earth, all physiological systems have been optimized for life in gravity.[142] This particularly is obvious for the vestibular system but also is true for the cardiovascular system, the sensorimotor system (which is responsible for movement coordination and control), and the system of bones and muscles. In fact, there is almost no physiological system in humans that has not been shaped by the specific gravity conditions on Earth. As a consequence, lacking this force in space induces several

---

[141] The lattermost option is, of course, extremely unlikely, given he is a fully trained astronaut with an advanced grasp of mathematics and astrophysics.

[142] Given the discussion here, this should be "life in a G-force environment of 1G".

physiological changes that call for complex adaptive processes in the human organism. (Kanas and Manzey 2008, 16)

Accordingly, tests of astronaut health on space stations make a significant contribution to our understanding of the effects of weightlessness on the human organism, and to the future health of astronauts on space missions. With overriding interests in such effects, one wants to test in actual conditions of weightlessness, as best one can, and not *apparent* weightlessness (much as a good baker wants a scale that gives the actual, not the apparent, weight of his ingredients). And so it is far from clear that, in this context, the "apparent" is apt or even intelligible, without calling to mind directive norms from other contexts of application which are not here salient. Moreover, given the Einsteinian interlude, it may be that the everyday directivities of "weight" are not as errant even in spacestation contexts as may seem: the equivalence principle sees Alfred on the orbiting ISS and Alfred at the Lagrangian point as equally weightless.

At any rate, Wilson elsewhere sounds liberal enough on these matters. For instance, he is perfectly willing to accept that there are other patchwork structures, such as that provided for the usage of "hardness", which inevitably exhibit multi-valuedness and yet could never usefully be flattened out in the way hoped for "weightlessness" (indeed, there have been failed attempts to do so for "hardness", as he discusses). Wilson later speaks of predicates that only *appear* to be "purpose independent" (2006, 455), but which actually are dependent, such as "hardness", "red", and (surprisingly) "sounds a low C tone". Purpose-dependence is a useful notion, somewhat already covered by talk on my end about context-dependence (recall my Chapter 3 discussion of the astronomer testing Mars's material for wetness; that it very much depends on his purposes as to what is correct to say). Likewise, in Wilson's terms, my point is twofold: (a) the content of "weighs x pounds" is purpose- and context-dependent (should those prove to be different qualities), and, (b), the content of most, if not all, terms is likewise sensitive to purpose and context.

The above discussion presses the case for (a): there are distinct Einsteinian, Newtonian, and everyday norms which license different applications of the "weight" predicate. The next claim, however, is prima facie radical, and it is not one which Wilson endorses. Thus, my argument for (b) in the next section will fully bring out the locus of disagreement between Wilson and me, especially on the relative contribution to conceptual content made by normative practices.

§4 Word/world alignment

In his opening chapter, Wilson flags up a difference between the sorts of predicates he is interested in and those entertained by "comparable literature of recent vintage", such as "dog" and "doorknob" (2006, 12). He has in mind a philosopher such as Fodor:

> I'm interested in such questions as: "What is the structure of the concept DOG?" . . . And my answer will be that, on the evidence available, it's reasonable to suppose that such mental representations *have no structure*; it's reasonable to suppose that they are atoms. (Fodor 1998, 22)

Views such as this are commonly paired with a thesis about content—say, the significance of the concept "dog" is exhausted by the attribute to which it refers (dogs), such that any satellite claims about dogs are external to the core conceptual content of "dog" (e.g., Wilson 2006, 60–61). What Wilson opposes here is not this view of conceptual content *simpliciter*, just its generalisation to every predicate. And so by focusing on predicates of this apparently simple variety, other literature, he claims, tacitly assumes that all concepts behave in the same way. Wilson will have no truck with that assumption: for him, predicative employments of "dog" and "hard" are importantly different. Any attempt to construe "hard", "red", or "rainbow" on the model of "dog" will likely have bought into the idea of semantic finality as earlier discussed.

Nonetheless, Wilson suggests that many concepts work in the way that he thinks "dog" does, where "predicate" simply aligns with worldly *attribute* (or "P"/φ for short). This is a point I cannot endorse. In the last section, we saw how it was not possible to flatten out the usage of "weighs x pounds" so that, across different patches of usage, it always aligns neatly with the same property. The problem was that there were simply too many normative concerns to separate out. In my view, such normative complexities—relative to purpose and context—are thoroughly pervasive, and they affect even the apparently simpler kind of predicate, such as "is a dog".

Now, to be clear, Wilson does not embrace the view of concepts he would find in Fodor— quite to the contrary, he insists it cannot adequately capture a number of more interesting concepts which experience "seasonal shifts" in their directivities (2006, 26); for those concepts, correct usage on different patches is guided by "the press of salient circumstance" local to those patches (2006, 43), not held firm by some invariant referential content (as he thinks "dog" probably is). The vocabulary subject to such seasonal variation is, in Wilson's view, normally that upon which we place great evaluative demands in some venture. For example, various industries apply "is hard" according to different standards, because they are tasked with producing a range of hard materials (2006, 100). In principle, then, for Wilson nothing exactly precludes our needing at some point to

place greater evaluative demands on "dog", such that its content would best be articulated not by a single attribute in the world to which it refers but by describing a patchwork of context- and purpose-sensitive usage. His main point here is simply that, in fact, we place relatively few demands on the large majority of descriptive or classificatory terms, such that appeal to "P"/φ alignment adequately articulates their conceptual content. As such, Wilson is recommending a softer kind of classicism about concepts; for him, semantic finality is true, up to a point, but it cannot be the whole story:

> my own position is quite clear: circumstances where a group of predicates sit in simple "is a dog"/*being a dog* relationships to the world are rather common, a fact that generally, but not always, represents a desirable semantic situation (almost certainly the predicate "is a dog" itself falls within this happy state). (Wilson 2006, 382)

Against this backdrop, I want to show that the application of deceptively simple concepts proves equally to be subject to the seasonal variation Wilson flags up for more demanding concepts. In other words, even the content of simple classificatory concepts is not exhausted by referential relationships with worldly attributes, and is more similar to concepts such as "hardness" than Wilson notices. We turn now to the work of Gilbert Ryle to make this case.

§4.1 Ryle's systematic ambiguity

Ryle distinguishes the ambiguity of so-called pun-words (such as "bank") from a more "systematic" variety (Ryle 1945, 215–17). He points out that we can easily resolve confusion stemming from pun-words by paraphrase or translation, and that the different significances of, say, "bank" *qua* financial institution and "bank" *qua* riverside land are normally made plain by the context of utterance.[143] On the other hand, *systematic* ambiguity, according to Ryle, is more pervasive and less visible than pun-word ambiguity:

> there is another sort of elasticity of signification which characterizes the use not of a few but of most or of all expressions and which is such that the paraphrases and translations of an expression with a certain elasticity of significance will normally have a precisely similar elasticity. (Ryle 1945, 215)

When Ryle says that paraphrases and translations of expressions carry with them the same elasticities as those with which the expressions begun, he is pointing to the internal character of

---

[143] Incidentally, these different senses of "bank" actually share an etymological root in a translation of "bench".

those abstracted expressions.[144] In other words, propositions and concepts are at least partly individuated by the common elasticities of significance they exhibit. In Ryle's terms, that is due to their "logical powers":

> [Any proposition] follows from some [others] as a consequence and it implies others. It is incompatible with some and merely compatible with others. It is evidence strengthening or weakening the probability of ulterior hypotheses. Further, for any logical powers possessed by a given proposition it is always possible to find or invent an indefinite range of other propositions which can be classed with it as having analogous logical powers or, as it is commonly put, as being of the same logical form. (Ryle 1945, 207)

To get a handle on Ryle's talk of logical powers, note its resemblance here to how Brandom discusses the inferential articulations of conceptual content.[145] Both philosophers would agree, for example, that a proposition such as "Amy is recuperating" presupposes "Amy has been ill"; further, both would take it that such a relation is underwritten by conceptual content. In Brandom's case, he will state that the content of the concept "illness", say, warrants this transition, and in turn that its content is supported at base by implicit "material" proprieties in practice. I am unsure whether Ryle would offer quite such a reductive account. In any case, Ryle is still happy to identify concepts with the families of propositions involving them, indefinitely many in number, employed in daily linguistic transactions (1945, 209).

Now, Ryle claims that all or most of our expressions undergo systematic ambiguity;[146] further, as discussed, expressions are *systematically* ambiguous because they exhibit unique sets of logical powers. These logical powers interrelate with those of other expressions, and carry over to paraphrases and translations of those expressions. Another way in which expressions exhibit this systematic kind of ambiguity, according to Ryle, is that they are employed in sentences with subtly different but related connotations: "The various ideas expressed by an expression in its different uses are intimately connected with each other" (1945, 215). To use Ryle's example, we say distinct

---

[144] I do not mean to invoke here Kaplan's unrelated notion of character, which takes contexts as input and outputs contents; in these terms, "I" has a non-fixed content: its content varies with the context of utterance. See Kaplan (1979, 83) for more details.

[145] Perhaps this is not surprising: Brandom's most salient inspiration is Sellars, in whose work approving references to Ryle are not uncommon.

[146] Kremer (2016) remarks that Ryle's term, borrowed from Susan Stebbing and possibly Russell, is unfortunate. He opts for "polysemy" as opposed to the pun-word "homonymy". I do think "ambiguity" is misleading insofar as it implies different *senses*, which seem far too coarse to indicate the much subtler differences Ryle has in mind.

but related things when we describe a man on a particular occasion as punctual compared to describing his arrival itself as punctual, or a whole class of people as such—the different employments of "punctual" within these several propositions "are in one way or another different inflections of the same root" (1945, 216–17).

Admittedly, this example is not especially revealing, but luckily in Chapter 3 we surveyed Ryle's early discussion of the different uses of "about", which demonstrates well the same kind of systematic ambiguity (Ryle 1933). In response to Williamson's attempt to discern what "aboutness" might amount to, we canvassed different ways to apply "about": for example, to the nominative subject of a sentence, to the object of a sentence, or to some topic in a conversation. But beyond these general pointers from Ryle, I showed how in conversation what sentences are about is heavily influenced by the salient concerns and purposes embedded in a context. In consequence, there are subtly different but connected uses of "about", but they do not carve out definitively different *senses* of "about", if the notion of a sense here might imply quite separate uses. Rather, as Ryle would later come to describe such linguistic phenomena, these different employments of "about" are inflections of the same root; they are systematically ambiguous.

In effect, then, Ryle's conception of systematic ambiguity invites us to reflect on the contextual application of expressions and, thereby, to notice the nuanced though systematic differences they enjoy. This is fully characteristic of how Ryle sought to analyse mental-conduct expressions such as "believes", "perceives", "thinks", "knows how", and "knows that" (see, e.g., Ryle 1949; on knowledge specifically, see Kremer 2016, 6). Additionally, his explicit examples include "exists" and the aforementioned "punctual". Julia Tanney has argued that "red" undergoes the same sort of elasticity, along with "reason", "intention", "understanding", and specific forms of expression such as "acting for a reason" (2013, 340ff).[147] Further, my own treatment of "bachelor" and "metre-stick" in previous chapters exhibits similar elasticities of significance. If right, Ryle and Tanney are pointing to a pervasive normativity which seeps through our language: the content of a concept is sensitive to what is *counted as* falling under that concept in the given circumstances— where circumstances, in turn, are strongly shaped by salient purposes.

§4.2 Two complaints

As such, the Rylean picture of conceptual content bears many affinities to Wilson's patchwork structure. However, despite some technical nuances, the glaring difference is in the absence of

---

[147] Tanney's discussion of "red" makes for an interesting contrast, including some overlaps, with Wilson's own on "redness" (Wilson 2006, 454–67).

normativity in Wilson's picture, in two respects. First, his purpose- and context-sensitive treatment of conceptual content does not extend to all or most concepts; it is more common, he maintains, for concepts to function in an unproblematic "P"/φ alignment manner (2006, 381–382; see also 60–61 on the "capacity-independence" of 'dog'). Second, even for those wandering concepts which Wilson treats as being supported by a patchwork structure, the role that norms play in contributing toward their content is considerably muted. For instance, as quoted much earlier, Wilson argues that whatever norms we learn as part of learning "is hard", physical obstacles encountered in working with materials are what actually push us far beyond linguistic rules.

With respect to the first complaint, I think we have reason enough to regard the content of such concepts as "dog" as still elastic. For scenarios are forthcoming in which "dog" is not properly applicable to the "attribute" with which it supposedly aligns even in the presence of, say, a member of the *Canis familiaris* species. That is, the predicate "is a dog" will not contribute to the content of all propositions in a uniform manner; it, too, is systematically ambiguous in Ryle's sense. For example, imagine that Superintendent Sharon, who controls the dog unit in the police station, has a pet poodle named Poppy. For whatever reason, Sharon must bring Poppy to work on one busy day. On that day, Sharon has had to send out all the dogs in her unit for cases before noon. Constable Bob, at noon, urgently needs a dog to accompany him on his own case. Bob impatiently says to Sharon, "I need a dog." But the Superintendent shoots down his request: "All the dogs are out, Constable." Bob grumbles, absent-mindedly looking at Poppy, and walks away.

How best to evaluate the truth of what Sharon says, in this context? Plausibly, she has not said anything false, and Bob indeed took her at her word. One might suggest that Sharon spoke elliptically for "police dog", and that Bob likewise understood this. Yet we are not compelled to read the scenario this way, for in *this* context, Bob needed a dog for a particular purpose: he needed a creature capable of performing certain tasks, for the likes of which poor Poppy has received no training. Poppy, so to speak, is not even in view for Bob or Sharon: what it is to *be* a dog in this circumstance is to be a creature of a particular constitution, ability, and certification. Perhaps here some will want to dig in their heels, and insist that Sharon's usage is nonetheless elliptical; I rather suspect that such insistence will be motivated by a felt need to stay firm to a theoretical view about what word-meaning *must* be, and how it must contribute to the truth of sentences. What is notable here, however, is that Wilson clearly has no such commitment, and in fact has at his disposal a ready explanation for this scenario: the predicate "is a dog" has dragged itself from the property *member of Canis familiaris*, say, to (a more specified) *creature of a particular constitution, ability, and certification*. Accordingly, in Wilson's terminology, the predicative use of "dog" is likely covered by a patchwork structure which is uneven (Wilson 2006, 323–24).

My response to the first complaint helps provide an answer to the second: that norms are only involved in a rather limited way for those wandering concepts. Wilson articulates this complaint by taking the "ordinary language school" as his foil:

> We agree that a term like "hardness" is governed by a complex schedule of localized controls in a manner that allows [metallurgists] to advance apparently contradictory claims in a coherent and informative way. But the ordinary language school also presumes that such communicative capacities can only be acquired as "rules" imbibed in the course of becoming competent in English. I, on the other hand, view the facade of hardness as gradually framing itself as an initially unspecialized term adapts itself, in a fairly predictable manner, to the individualized aspects of the materials around us. (Wilson 2006, 343–44)

One can see what Wilson is driving at here. From one patch of usage to another, metallurgists may make seemingly contradictory statements about the hardness values of a piece of metal—its hardness is given by Test A, but in the next moment, they say that Test B reveals its "true" hardness, and so on. Their ability to transition smoothly between such multi-valued judgements is ensured, Wilson thinks, by the control structures governing "hardness" they implicitly adopt in the course of investigation.[148] That is, they learn that for some given metal needed for such-and-such a purpose, its hardness is best tested by scraping its surface; meanwhile, its hardness for some other purpose is better tested by compressing, and so forth. And such strategic operational skill, encapsulated in localised controls, is surely not "guided" by linguistic rules. Rather, Wilson concludes, such rules are "overcome by the impertinent particularities of metal, ceramic and plastic" (2006, 344).

On the face of it, then, normativity takes second place in determining the conceptual content of "hardness", for linguistic rules can prove inadequate when predicates wander into regions with more exacting and novel demands, as commonly found in materials science. However, to my ears, Wilson's localised controls are normative through-and-through. True, one may not articulate the recommendations of hardness-testing so explicitly *as* rules, even in learning them through practice. But one will still be adapting usage to normative concerns—what counts as hardness in this context differs from that context—and one can articulate them as rules to guide others if need be.

Moreover, it seems like the true locus of disagreement between Wilson and his (imagined) ordinary language critic concerns the driver of normative change, as it were. As we see in the next section, Wilson describes linguistic evolution in terms of an interfacial compromise between our

---

[148] "Control structure" is an earlier term used by Wilson (1988) to depict a similar idea to his patchwork atlases.

linguistic abilities or tools and the physical traits themselves (the particularities of metal, say). Wilson imagines the ordinary language proponent arguing that our ability to adapt seemingly contradictory usage to specific domains is down to the " 'rules' imbibed in the course of becoming competent in English" (2006, 344). I am unsure who holds exactly that kind of view; as far as I am concerned, this is a simplistic and unhelpful picture of the normativity of language.

Indeed, surely what needs emphasising here in addition to Wilson's interfacial talk, as a prompt for normative change, are two more points. First, that the conceptual content of our classificatory terms is partly determined by purposes relative to conceptual context. For what purpose do we want a given piece of metal to be hard (as armour, penetration, load-bearing, and so on)? I contend that it is only with respect to those purposes that it makes sense to determine the content of "hard" in a given context. (This point will play out significantly in the next section, too.) The second point is that Wilson's suggested localised control systems, the patchwork facades we operate with, constitute a form of practical know-how. If that is true, as I argue more plainly in §5, then in a sense one's adaptation of "hard" to new investigative purpose, given the response of physical traits, is of course conceptual and normative.

Finally, consider Sharon and Bob again. How are they able to apply and respond to "dog" in the manner described in my scenario? Imagine that after walking away from Sharon, Bob utters to a colleague, "I can't stand Sharon's dog—silly looking thing." Bob can say this and yet *still* be committed to the claim earlier heard, that there are no dogs in the station. Prima facie it seems like he is again entertaining a contradiction, but this is an especially uncharitable view of Bob. Rather, it seems to me that he has picked up on the salient and implicit norms of the different scenarios: in the first, what counts as a dog is not what counts as a dog in the second, given different concerns and purposes.[149] Wilson himself comes close to such an account when he explains how we can handle multi-valuedness in a patchwork structure; recall that for "weighs x pounds" there are so-called prolongation rules which we follow when jumping from a patch covering a terrestrial domain to a space-shuttle domain, such that we are not genuinely contradicting ourselves when saying that Alfred weighs 0lbs and 160lbs. Accordingly, given his appeals to what seem to me clearly normative concerns, it is hard to see how Wilson's attempt to demote the normative in his picture of conceptual content could succeed.

Perhaps a large part of the problem here is that Wilson presumes any fairly robust appeal to normativity, in explaining the content of a concept, must be of a semantic finality stripe. That is, Wilson only seems to think one says "The proper use of concept 'C' is…" when one has in mind

---

[149] Cf. Tanney's account of the shifting criteria for "love" (2013, 347).

a precisely delimited usage, such as a norm taught as part of learning the language. But Wilson's own picture is much more forgiving, and it can be read in normative terms too. It may be that purposive empirical investigation into some subject prompts us to change course in how we apply a concept, yet still that change is expressible against a backdrop of prior normative employments, and itself has normative potential. Accordingly, there is no escaping the fact that conceptual content is normative.

§5 The looming veil

In the previous sections, we sought to reject some of Wilson's offerings. We could not accept that Wilson's pointing to interactions between control structures and physical traits fully gets a concern with "rules" out the picture. Nonetheless, there is still much to keep: a picture more detailed than any offered elsewhere of specifically empirical predicate usage; a more connected picture thus drawn of different usages, strung together as patches (and no longer so disparate, as were, seemingly, my "pockets" of Chapter 3). So much, it seems, we can retain while arguing substantive points against Wilson on the score of normativity. However, a new tension surfaces when taking into account an enduring strain of Wilson's argument.

Throughout *Wandering Significance*, Wilson recurrently recommends a diet low in certain Kantian (especially *neo*-Kantian) insight. He is keen to warn against a particular neo-Kantian thought concerning objectivity:

> In rough terms, the general claim [to be avoided] is that our naïve conception of "objective" concepts as *correspondent* to real world attributes is incoherent; that every viable concept must inherently involve the constructive agencies of our own minds in some irrevocable way. (Wilson 2006, 77)

On the face of it, this sounds like an accurate depiction of Kant's view. Recall, in Chapter 1, our discussion of Kant's thought that objects conform to reason. For Kant, our a priori intuitions of space and time, being the forms of perception, do not originate from perception itself but from us. Further, via the faculty of understanding, we employ a priori concepts (categories[150]) to synthesise what sensibility throws our way, a manifold of representations. Thus, by Kant's reckoning, the subject plays an active role in organising what it is that she experiences. This is Kant's so-called Copernican turn: to argue that the objects of knowledge and experience (setting

---

[150] Specifically, the categories of universality and necessity—see the Transcendental B-Deduction, especially §§16-20 (Kant 1787).

aside the supersensible) are not external to our understanding, but rather are "objects-for-us". So if "real" in Wilson's "real world attributes" is supposed to describe objects which have not conformed to reason—presumably, things-in-themselves—then, indeed, Wilson's implicit correspondence claim would fall foul of Kant. The good Kantian line by contrast is that our concepts of objects are objectively valid[151] (and, I imagine, "correspondent") so long as we do not understand "the world" to comprise things-in-themselves, a region in principle alien to the faculty of understanding.

Now, the reader may have noticed that throughout this thesis there is a strong Kantian bent; I am sympathetic to Kant's attacks on traditional empiricism, which may well carry over to the contemporary naturalism of which Williamson is a fair representative.[152] Moreover, Sellars's fears of the Given, and McDowell's attempted solution to the problem posed by it (both of which I endorse), are much less compelling if one does not take seriously the view of Kant just discussed. Thus, if Wilson is right to warn the reader away from this basic Kantian insight, it spells trouble.

§5.1 Kantian tension

Wilson's resistance to neo-Kantian views finds its clearest expression in his warning not to succumb to the "veil of predication" (Wilson 2006, 390–401). The veil is described loosely by Wilson; indeed, it is hard to get a clear view in our sights of exactly the phenomenon he targets. Worse, he accuses wide-ranging philosophical positions and philosophers as having succumbed to it: Fregean views of sense, especially "mode of presentation" talk (2006, 303, 396); Russell's claim that we know universals only through description (2006, 394); Quine and other pragmatists in certain of their holistic moods (2006, 261, 279), as well as Kuhn (2006, 82); of course, neo-Kantian work, Heideggerian views, communitarian readings of Wittgenstein, and Sellars too (2006, 83–84); finally, and most explicitly, Nelson Goodman (2006, 397–98).

According to Wilson, just as traditional empiricists found the world behind their postulated sense-data essentially unreachable (behind a "veil of perception"), so too do those who fumble under the veil of predication (2006, 81), although this time the world is hidden behind concepts, not sense-data. The motivation for such a view, apparently, is an overzealous rejection of

---

[151] Note, there is some terminological nuance that I am liberally brushing past here. For one thing, an object in this discourse, for Kant, is more an objective feature of reality than what may colloquially be interpreted by "object". See Pereboom (2014, §1.2-1.3)

[152] See Redding (2007) for a discussion of how Hegel's critique of Kant's "supersensible" realm also may apply to modern scientific naturalism.

word/world (or "P"/∅) correlation, in which (as we have recently discussed) a worldly property, universal, or sense-datum is thought to underpin the conceptual content of a term, much as the property *Canis familiaris* supposedly supports "dog". Those keen to avoid a simplistic correlation picture—whom Wilson dubs "anti-correlationists"—are inspired by Kant to maintain that "every viable concept must inherently involve the constructive agencies of our own minds in some irrevocable way" (2006, 77). The anti-correlationist thought, as I take it, is that any given conceptual content cannot entirely consist of that to which it refers; the predicates "is a dog" and "is round" do not latch onto real-world dogs or round things (or sets of round things) and thereby attain their content. Rather, human subjectivity plays some intervening role in forming these conceptions. There can be many glosses on such a view; many different rationales for the general claim. But among the worst, thinks Wilson, are those which postulate as equally available to us several incompatible theories describing the world, each of which on its own can be said to describe the world correctly (2006, 81).

Indeed, Wilson's greatest fear with respect to anti-correlationist talk is that it prevents us from appreciating the objectivity of our language—its "patterns of genuine correlation" (2006, 80), both simple word/world correlation and more complex patchwork types. The veil of predication, he claims, "bars us from ever determining whether the concepts we employ genuinely match the true traits of the world or not" (2006, 82). A more mature picture, by Wilson's reckoning, acknowledges the human contribution to linguistic evolution, but not at the expense of objectivity: e.g., we learn to navigate patchwork facades with special techniques in order not to assert multi-valued claims, and in so doing we employ our idiosyncrasies while retaining objective information about the world (2006, 400). Accordingly, Wilson's alternative picture stresses that "interfacial factors" shape our linguistic usage in unnoticed ways: our linguistic or "representational tools" adapt to "physical circumstance" (2006, 288). It is at this interface between our capacities and the world that patchwork facades form and allow classificatory conceptual content to correlate with physical traits. Wilson's worry about the veil of predication, then, is that it seemingly forbids by fiat any kind of conceptual correlation with the world. Thus,

> philosophical sermons to the effect that it is inherently incoherent to segregate the subjectively based aspects of linguistic shaping from their more objective counterparts do not represent music to my ears. (Wilson 2006, 83)

One such sermon to which Wilson takes exception is Goodman's:

> If I ask about the world, you can offer to tell me how it is under one or more frames of reference; but if I insist that you tell me how it is apart from all frames, what can you

say? We are confined to ways of describing whatever is described. Our universe, so to speak, consists of these ways rather than of a world or of worlds. (Goodman 1978, 2–3)

Wilson finds it problematic that Goodman thinks we cannot rightly point to one "world" or worldview as the correct one, and thus to ratify some set of objective facts. I do not wish to enter this specific debate—not least because the idea of a plurality of worlds is prima facie technical and confusing—but it strikes me that we *ought* to offer some "sermon" to the effect that the meaning of even our classificatory language is everywhere infused with our practical, normative, conceptual schemes. Indeed, I think that is the lesson to have extracted from the McDowell of *Mind and World*, who, strongly influenced as he is by (Strawson's) Kant (McDowell 1994, viii), seemingly makes an apt target for Wilson.

§5.2 The Given

To remind, McDowell is struck by an epistemological problem which should give us pause: what provides our empirical beliefs (i.e., beliefs about the world) with rational support? Are they justified by "the world itself", or maybe by some parts of that world? In the sense of "to justify" in which I say "Judge so-and-so's decision is justified by Article x and cases y and z",[153] what justifies my belief that there is a tree in front of me, for example? There is an easy temptation to say here, simply, "the tree"—after all, were there no tree, my claim should not be true. But this is not the sense of justification to which we just appealed. When I say that there is a tree in front of me, I am using language to describe how things are. If I am asked "Why?" in response to my description, I am put in a position to offer a reason; I must show my specifically *rational* entitlement to the belief. For example, think of how the fact that a tree is before me itself provides reason to believe further claims, such as that I am not in a sealed basement, or that what lies before me is an organism. Accordingly, there are rational relationships between claims here, and I am required to traverse these pathways to offer a reason for my empirical claim.

The problem with the answer "the tree" or "the tree itself" is that this part of the world as such offers nothing in the form of a reason itself. The tree is rationally mute, as it were. The presence of the tree is essential in this case for the accuracy of the belief, but the tree is not that to which the belief is *answerable*. In order for a listener to be satisfied that a speaker understands her words, she oftentimes will need the speaker to offer reasons for his claims; in this sense of "reason", the

---

[153] I use a legal scenario here as an example where claims and decisions demand rational answers within normative frameworks.

tree itself is not a reason. Not in the way that "because there is a tree right in front of me" is such a suitable reason; this latter response shows a rational entitlement. In short, anything which lacks this conceptual togetherness, as it were, is not the right category of thing to provide a reason for a belief. And anything which purports to provide such a reason for belief while being non-conceptual is thus a form of the Given.

One temptation in response to this failure of the Given is to forgo any attempt to make the world answerable to our beliefs. In the vein of McDowell's Davidson, one might cite the world's causal powers as providing at least some external constraint. Yet, as McDowell argues, this ignores a key insight: that we, in the world, are normally rationally entitled to our beliefs about how things are, and a cause is as rationally inert as the Given itself (or as the tree itself in my example).

As we are by now familiar with, McDowell's solution to the problem of how to undergird our beliefs is to show that our experience of the world is in a sense conceptual: it is imbued with conceptual capacities. One could not, for example, discern a tree *as* a tree without possessing a concept of a tree. And so, the world we experience is one which is already conceptually articulated; we are at any given moment able to perceive and describe it adequately only in virtue of our conceptual abilities. This means that even in our passive perception of the world, what we see is already conceptualised. In McDowell's terminology, those conceptual capacities are always operative, even if not being exercised or actualised; accordingly, for McDowell, the content of perceptual experience itself is always in a sense conceptual. Further, if one is to classify, in one's experience, an object as "red", "large", "bright", and so on, one is bringing to bear conceptual capacities that one must be able to actualise outside this classificatory experience:

> Quite generally, the capacities that are drawn on in experience are recognizable as conceptual only against the background of the fact that someone who has them is responsive to rational relations, which link the contents of judgements of experience with other judgeable contents. (McDowell 1994, 11–12)

Now, McDowell's account of perceptual experience sounds to my ear prima facie similar to the views which bother Wilson. For Wilson thinks much of our language correlates with objective traits in the world, either in simple "P"/ϕ fashion or via patchwork facades (or other, more complex structures), and yet McDowell's solution to the problem of the Given is to say that our experience of the world is always in some sense mediated by concepts. Perhaps, however, "mediated" is tendentious, for the world's "deliverances" are, in McDowell's view, conceptually structured. Indeed, this last thought helps untangle the present knot.

The relevant question to ask here is thus the following: does McDowell's account of perceptual experience sink into the kind of idealism which refuses objectivity? It is only by subscribing to

such an idealism that my own discussion of conceptual truth and normativity could find itself in significant tension with Wilson's position.

§5.3 Whose objectivity?

For McDowell, the objective world is not irretrievably lost in its being conceptual. Yes, when I take my experience of a red ball at face value and report "the ball is red", I am exercising the conceptual capacities already implicit in my experience of the ball *as* red. However, it is a fact of the world too, since *that the ball is red* is a perceptible fact. More generally, *that things are thus and so* is conceptual content but it "is also, if one is not misled, an aspect of the layout of the world: it is how things are" (McDowell 1994, 26). Not only, indeed, is this a way to save objectivity; as far as McDowell is concerned, it is the *only* way. For the alternative, he argues, is either to postulate the Given, which is mythical, or to exculpate our judgements from answerability to the world (1994, 26–27).

How does this tally with Wilson's conception of objectivity; is McDowell redeemed on that score? We can only tell that once we have Wilson's conception of objectivity firmly in mind. For Wilson, objectivity is "concerned with the manner in which language finds *correlated underpinnings* within the world before us" (2006, 78–79). Thus, the ball really is red because it exhibits physical traits correlated with "red"—the attributes we "encode" in such judgements as "the ball is red". Further, as Wilson shows, "red" is a classificatory predicate which we apply with great strategic ingenuity (unlike, as we have seen, he imagines for "dog"), of which we are mostly unaware. For example, we tacitly make vital discriminations of texture and incorporate aspects of the surrounding environment in even humble judgements of a given object's colour.[154] Having followed strategic perceptual procedures home to the patchwork structures of colour, surface, and texture predicates, we are well placed to judge that the ball is red. However, Wilson claims that this human contribution to the colour report does not hopelessly infuse it or its constituent concepts with subjectivity:

> In my story, humanly idiosyncratic factors can shape the *personality* of a predicate like "red" greatly without compromising its capacities as a carrier of objective information. Quite commonly, the most important considerations behind the formation of complex facades have little to do with issues of subjectivity (in the sense of human psychology)

---

[154] See Wilson (2006, §7(x)) for a comprehensive tour of colour vocabulary employment in everyday and technical settings.

at all, but are determined by the measurement and inferential tools readily available at a macroscopic level. (Wilson 2006, 400)

Wilson's story here is that our usage of "red" is shaped by a compromise between our capacities to infer and measure, on the one hand, and the physical traits that we want to record, investigate, and discuss, on the other.

I presume, in addition, there is a hidden pragmatist element at this interface (Wilson classifies himself as a kind of pragmatist, after all), because the practical purpose of using inferring and measuring "tools" is always lurking in the background. Indeed, it must be with respect to such purposes that "red" can be said to encode objective data; after all, different patches of usage— different contextual employments—are partly demarcated by difference in *point*. For example, it matters for a different reason why the ball is "red" on a snowy football pitch to when the ball has rolled into a patch of tomatoes. Accordingly, here is another element of human idiosyncrasy in which our language is enmeshed, while it hardly seems to prevent us from accurately describing red items in the world. In fact, not only does purposiveness fail to hinder objective description, it is arguably an enabling condition—think back to the purposes I ascribed to the Martian astronomer in Chapter 3; without any purposes what significance has his enquiry and his use of "dry"?[155]

Thus, Wilson rescues his own picture from untenable subjectivism by appealing to what seems to me a pragmatist criterion of objectivity. That is, worldly data is objective insofar as it is recorded in ways that allow us to infer, predict, explain, understand, investigate, and pursue our exploratory endeavours. *That* criterion is surely compatible with a Kantian conception of objects, according to which in order for any "physical traits" to be hypothesised or investigated, we must apply concepts to them. More significantly, this conciliatory harmony between the pragmatist and Kantian thought should carry over without problem to McDowell's philosophy.

Indeed, McDowell's talk about the background conditions prerequisite for experience of colour are compatible with Wilson's discussion of our strategic employment of colour predicates. (Moreover, it seems that McDowell understates how extensive those background conditions are, given Wilson's survey.[156]) Importantly, McDowell regards these tacit abilities as conceptual; a colour-observer is "equipped with such things as the concept of visible surfaces of objects, and the concept of suitable conditions for telling what something's colour is by looking at it" (1994, 30). And clearly, the kind of concept-possession that McDowell has in mind is more a form of

---

[155] See Chapter 3, §2.2 for more discussion.

[156] See the footnote before last.

know-how than a form of know-that: one might, upon reflection, be able to articulate guiding norms such as "Look at the ball in uncovered daylight to see its colour", but one may just as easily never think such and yet still follow the norm by living in accordance with it, so to speak. Likewise, such a tacit, guiding norm is not essentially concept-forming; at least, we would not expect someone to articulate the norm as a way of demonstrating their grasp of the concept "red" (say). Nonetheless, the practical skill so employed on given occasions would manifest a kind of knowledge about colours, objects, surfaces, and so on. It is a small step from here to stating that "red's" guiding directivities (to use Wilson's terms) are in the same sense conceptual: we follow such procedures in order to ascertain, confirm, and pronounce the colours of objects (among other purposes).

At this final juncture, it is possible, I think, to tie neatly together the different strains of thought here considered: Kant's conception of objectivity, McDowell's conceptual content of experience, and Wilson's pragmatist gloss on objectivity (as I read it).

Consider Pereboom's reconstruction of Kant's Transcendental B-Deduction, in which Kant hopes to show, *inter alia*, how Hume's associationist psychological theory could not account for the representation of objects. Suppose one looks at a house from the front, and then the side, rear, and other side. Chronologically, one's sense-impressions of the front, rear, and sides are sequential, yet one experiences these phenomenal parts of the house as simultaneous, not strung together from different temporal impressions. Compare that scenario with watching a boat sail down a river: one's sequential sense-impressions convey in this case the boat's co-ordinate changes as temporally successive. A psychological theory which admits of no other source of knowledge than sense-impressions struggles to tell these cases apart; Hume would need an experiential criterion of simultaneity which seems unavailable for him. In Kant's case, however, he can appeal to the application of the categories *Understanding* and *Necessity*. Pereboom represents the import of those categories in the following principle:

> (U-N) Necessarily, if empirical conditions are normal, any human experience of the parts of the house is an experience of these parts as objectively simultaneous.[157]

Presently, it is not important whether there are issues with this specific formulation (for example, what its logical and modal aspects amount to). Rather, just notice that (U-N) seems a realistic articulation of conceptual know-how; the principle expresses the kind of conceptual capacity which lay operative in our experience of the world. That is, (U-N) stands shoulder to shoulder both with an expression of the capacities Wilson attributes to employments of "hard" and "red",

---

[157] See Pereboom (2014, §1.3) for further discussion of this principle.

and such of McDowell's capacities as might be expressed by "Look at the ball in uncovered daylight to see its colour". I contend that all such know-how forms a large part of what one must learn in order to participate in normative, linguistic practices.

Thus, this blend of ideas offers us a small glimpse of how diverse and subtle our conceptual capacities are, along with an appreciation of what conceptual content is.

Conclusion

This thesis addresses a question of contemporary philosophical relevance: How should we characterise the methodology of philosophy?

A common response, exemplified at the beginning of this thesis by A. J. Ayer (1956), is to say that philosophy investigates the meaning of words, though not the words themselves. Ayer describes this process as one where we look at an application of words with which we are familiar and see how it fares in real or imaginary cases; subsequently, we then question "whether there is anything that we should be prepared to count as an exception to the suggested rule [in such cases]" (1956, 28). There is a clear sense, then, in which for Ayer philosophy is a conceptual discipline: we arrive at truths principally through reasoning with and investigating our concepts in a multitude of possible scenarios. That is more a claim about method than it is about subject-matter; for Ayer also claims that such enquiry issues ultimately in *facts* (1956, 29).

Having worked through Williamson's critique of what he calls "conceptual" or "linguistic" philosophy, we can, surprisingly, see a similarity with Ayer here. Although Williamson argues that "few philosophical questions are conceptual questions in any distinctive sense" (2007, 3), one reason he thinks this is because he regards any putative conceptual truths as being true in the same "sense" as empirical truths. That is, he thinks "Adults are grown-ups" is true in the same, factual sense of "true" as "Adults are disillusioned"—the latter, if true, is just less obvious. But, as I have shown, things cannot be so (and partly because the idea of conceptual truth is not the idea of a different *sense* of "true").

The most critical point of this thesis is that the kind of equivalence which both Ayer and Williamson (along with Boghossian) draw is fatally wrong: facts about what is the case and norms enjoining how to describe what is the case are logically independent. What now functions as a description may subsequently express a norm, but on one and the same occasion of its utterance, no sentence can express a rule and describe what is the case. Consequently, just as token uses of "Adults are grown-ups" are typically normative (i.e., they license transitions from, say, "Laura is an adult" to "Laura is a grown-up"), so are the truths over which philosophy puzzles: "Nothing is red and green all over", "Red is a colour", "Only I can feel my pain", "Knowing entails believing", "The mind is not a substance", "Meaning is use".

And if one objects to characterising philosophical propositions as normative because one can conceive scenarios in which they are not true, then one has missed the point: what a sentence amounts to will be tied to its circumstances, within which, if anything is sayable at all, then there are norms governing what is intelligible to say. Thus, one may (and should) point out that

sometimes it is perfectly legitimate to say that one knows that *p* while not believing that *p*. For example, "I don't believe it happened; I *know* it did". The reason this is legitimate is because the uses of "believes that" and "knows that" subtly vary in significance, as circumstances and purposes change. For instance, sometimes, "A believes that…" expresses a lack of confidence, other times acceptance, or a qualified estimate;[158] such uses are often incompatible in circumstances where a higher degree of assurance is sought. In each case, these uses are legitimated by norms, the kind of norms which philosophy expresses as analytic, conceptual truths. Far from undermining philosophy's claim to conceptual investigation, then, such an objection reinforces it.

At the outset of this thesis, I sought to engage with figures such as Boghossian and Williamson by agreeing to a central tenet of both their separate metaphilosophical views: that analytic truth and conceptual truth are the same thing. However, I qualified this agreement by insisting there is no one concept of analytic truth; instead, there are numerous conceptions. Chapter 1 was motivated to show how a criticism of one conception need not of itself present a danger to another. Another purpose was to demonstrate the considerable heritage that the analytic-synthetic distinction enjoys. To be sure, in light of that ancestry, there is reason to avoid describing what I call conceptual truth as analytic truth (Baker and Hacker 2009, 261–62; Hacker 2009, 343) or else to regard "analytic" as mostly though not entirely overlapping with "conceptual" (Schroeder 2009, 102–8). For example, unlike, say, Frege, I am not associating a truth's being analytic with its being arrived at by some process of analysis. Indeed, to my mind, it becomes inevitable that we segregate analysis (*qua* decompositional, transformational, or regressive) from analyticity as soon as we recognise the normative function of analytic truths in practice.

Moreover, that same realisation encourages a way of rethinking, rather than discarding, analyticity: since analytic sentences are the ones that express norms, why not defend the analytic-synthetic distinction afresh as one that distinguishes norms from descriptions? Such a defence, to my mind, is especially advantageous in a contemporary setting—analyticity has taken centre stage in metaphilosophical discussions, and yet disputants too regularly do not appreciate the role of normativity. Chapters 2 and 3 of this thesis show just how this lack of appreciation leads to two unsettling philosophical outlooks, in the work of Boghossian and Williamson.

Like myself, McDowell also hints, as we saw in Chapter 4, at the possibility of reworking analyticity to good effect. My attempt to do so in this thesis coheres with McDowell's explanation of how we can use our experience of the world to justify what we say about it. In the slogan I eventually opt for, this means that analytic truth is a precondition of synthetic truth. We learn a

---

[158] See Hanfling (2000, 101–5) for numerous interesting examples.

language by being embedded in normative practices, which we can articulate in the form of analytic truths. A handy way of putting the point is to say that our concepts are constituted by a network of normative, rational relations; to grasp a concept is then to know one's way around relevant portions of that network.

It is tempting, but unhelpful, to gloss this point as saying that conceptual content is exhausted by the inferential links of such a network. Tempting, because, as Brandom argues, anyone who understands a proposition must be able to specify what follows from it, or what entitles them to it—someone who asserts "The pitch is green" is committed to, say, "The pitch is coloured". Further, according to Brandom, so long as a person can make these inferences, and so long as she can reliably respond to her environment with statements such as "That's green", she is to be counted as knowing that the pitch is green. McDowell shows us, however, that such a subject is fully justified in her claim that the pitch is green when she sees *that the pitch is green.* True, she needs to know which statements follow from which others in order to be credited with having the concepts involved in the judgement, but what provides her entitlement to the judgement is the conceptually mediated experience itself—not her ability to reel off a list of true inferences. Accordingly, having availed ourselves of the idea that analytic truth preconditions synthetic truth, modelling perceptual justification on reliable responses to the environment (as Brandom does) is an unnecessary and unhelpful step. For if we have had the relevant normative training, we encounter the pitch's being green in experience.

Finally, in Chapter 5, we consider prima facie objections to my approach posed by Wilson. Wilson thinks that as language-speakers we commonly have a poor grasp of our own concepts, and this is partly because the norms which govern their application are, he suggests, sometimes a step behind Nature itself. Focusing predominantly on concepts from applied physics and mathematics, Wilson charts the ways in which predicates "wander" from signifying one property in one circumstance to another across the border. This shifting significance is something Wilson alleges we frequently do not notice, and which our norms have not yet accounted for. But the problem for my thesis is largely illusory; Wilson proves more ally than foe. For what Wilson shows, in my view, is not that norms fail to determine conceptual content; to the contrary, he shows that in technical disciplines the new methods of description which participants take up are conditioned by norms which extend application, or which show how a new application is relatable to an old application. For instance, by looking at the metallurgists' varying uses of "is hard", Wilson shows that the concept of hardness does not—and could not—admit of a uniform "flattening out", such that "hardness" could be defined as some specific kind of physical property.

What transpires in this final chapter, then, is that Wilson (somewhat like Williamson) turns away from talk of norms because he has in mind a rigid conception of them, where they fail to adapt to new circumstances. This view of normativity is representative more of a human slowness to articulate the norms they follow than anything like a slowness of norms in themselves. The way Wilson illuminates, effectively, a more fine-grained and interesting picture of normativity in practice is by use of so-called patchwork structures; these can be used to show how we cope, normatively, with shifts in domain and jumps between them. What is curious, however, is that Wilson is still wedded to a view on which plenty of predicates do *not* function in this way, but rather pick out properties across domains, in "P"/φ fashion. Wilson's paradigmatic example is "dog", which he sees as tethered to a species property. But, as I show in Chapter 5, even this term is still subject to elasticities of significance—and not of a straightforward "pun"-style ambiguity— according as contexts of employment vary. In each case, this semantic nuance is to be described as enabled by normativity, not as some flouting of it. That is, a seemingly errant claim such as "There are no dogs here", when a poodle is plainly in view, can be perfectly true and yet betoken no misunderstanding so long as there is a norm which does not *count* such a creature as a dog in the circumstances. In sum, then, Chapter 5 presents us with an opportunity to show how pervasive normativity is; in the technical and in the everyday.

Meanwhile, the overarching lesson of this thesis is that philosophy must take account of linguistic norms if it is to dissolve problems or to elucidate phenomena. Accordingly, we have to construe philosophy as a conceptual investigation, where this means that philosophy investigates, clarifies, untangles, and arranges the norms governing the concepts involved in characteristic statements of its problems. We can make this point by saying that philosophy is in the business of articulating conceptual truth. But too frequently that claim is assimilated to the thought that philosophy is about analytic truth, and analytic truth is disregarded. Let us go straight to the heart of the dispute, then, and speak in the terms of the disputants. The statements of philosophy are indeed analytic truths, because on closer inspection analytic truths are normative; this explains why they are not—for lack of a better phrase—"made true" by the world: they are not in the business of describing anything. A failure to recognise as much seriously undermines contemporary debates on the methodology of philosophy.

Bibliography

Allison, Henry E. 2008. *Custom and Reason in Hume: A Kantian Reading of the First Book of the Treatise*. Oxford: Oxford University Press.

Anderson, R. Lanier. 2015. *The Poverty of Conceptual Truth*. Oxford: Oxford University Press.

Austin, J. L. 1961. 'A Plea for Excuses'. In *Philosophical Papers*, edited by J. O. Urmson and G. J. Warnock, 123–52. Oxford: Clarendon Press.

Awodey, Steve. 2007. 'Carnap's Quest for Analyticity: The Studies in Semantics'. In *The Cambridge Companion to Carnap*, edited by Michael Friedman and Richard Creath, 226–47. Cambridge Companions to Philosophy. Cambridge: Cambridge University Press. https://doi.org/10.1017/CCOL9780521840156.011.

Awodey, Steve, and A. W. Carus. 2007. 'Carnap's Dream: Gödel, Wittgenstein, and Logical, Syntax'. *Synthese* 159 (1): 23–45. https://doi.org/10.1007/s11229-006-9066-4.

Ayer, A. J. 1936. *Language, Truth and Logic*. London: Penguin.

———. 1956. *The Problem of Knowledge*. Middlesex, UK: Penguin Books.

Baker, G. P., and P. M. S. Hacker. 2005. *Wittgenstein: Understanding and Meaning: Volume 1 of an Analytical Commentary on the Philosophical Investigations, Part II: Exegesis §§1-184*. Oxford: Wiley-Blackwell.

———. 2009. *Wittgenstein: Rules, Grammar and Necessity: Volume 2 of an Analytical Commentary on the Philosophical Investigations, Essays and Exegesis of §§185-242*. Oxford: Wiley-Blackwell.

Bar-Elli, Gilead. 2010. 'Analyticity and Justification in Frege'. *Erkenntnis* 73 (2): 165–84.

Baz, Avner. 2012a. *When Words Are Called For*. Cambridge, MA: Harvard University Press.

———. 2012b. 'Must Philosopherss Rely On Intuitions?' *The Journal of Philosophy* 109 (4): 316–37. https://doi.org/10.5840/jphil2012109412.

Bealer, George. 1987. 'The Philosophical Limits of Scientific Essentialism'. *Philosophical Perspectives* 1: 289–365. https://doi.org/10.2307/2214149.

Beaney, Michael. 2016. 'Analysis'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Summer 2016. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/sum2016/entries/analysis/.

Beck, Lewis White. 1978. 'Analytic and Synthetic Judgements before Kant'. In *Essays on Kant and Hume*, 80–100. West Hanover, MA: Yale University Press.

Boghossian, Paul Artin. 1996. 'Analyticity Reconsidered'. *Noûs* 30 (3): 360–91. https://doi.org/10.2307/2216275.

———. 1997. 'Analyticity'. In *A Companion to the Philosophy of Language*, edited by Bob Hale and Crispin Wright, 331–68. Oxford: Blackwell.

———. 2003. 'Epistemic Analyticity: A Defense'. *Grazer Philosophische Studien* 66 (1): 13–35. https://doi.org/10.1163/18756735-90000810.

———. 2011. 'Williamson on the A Priori and the Analytic'. *Philosophy and Phenomenological Research* 82 (2): 488–97. https://doi.org/10.1111/j.1933-1592.2010.00395.x.

———. 2012. 'Inferentialism and the Epistemology of Logic: Reflections on Casalegno and Williamson'. *Dialectica* 66 (2): 221–36. https://doi.org/10.1111/j.1746-8361.2012.01303.x.

BonJour, Laurence. 1997. *In Defense of Pure Reason: A Rationalist Account of a Priori Justification*. Cambridge, UK: Cambridge University Press.

Brandom, Robert. 1994. *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge, MA: Harvard University Press.

———. 1995. 'Perception and Rational Constraint: McDowell's "Mind and World"'. *Philosophical Issues* 7: 241–59.

———. 2001. *Articulating Reasons*. Cambridge, MA: Harvard University Press.

———. 2002. 'The Centrality of Sellars's Two-Ply Account of Observation'. In *Tales of the Mighty Dead: Historical Essays in the Metaphysics of Intentionality*, 348–67. Cambridge, UK: Cambridge University Press.

Burge, Tyler. 2000. 'Frege on Apriority'. In *Truth, Thought, Reason: Essays on Frege*, 356–87. (2005). Oxford: Oxford University Press.

———. 2003a. 'Postscript to "Frege on Apriority"'. In *Truth, Thought, Reason: Essays on Frege*, 388–89. (2005). Oxford: Oxford University Press.

———. 2003b. 'Logic and Analyticity'. *Grazer Philosophische Studien* 66 (1): 199–249. https://doi.org/10.1163/18756735-90000817.

Burnham, Douglas. 2017. 'Leibniz, Gottfried: Metaphysics'. Edited by James Fieser and Bradley Dowden. *Internet Encyclopedia of Philosophy, ISSN 2161-0002*. http://www.iep.utm.edu/leib-met/.

Cappelen, Herman. 2012. *Philosophy without Intuitions*. Oxford: Oxford University Press.

Cappelen, Herman, and Ernie Lepore. 2005. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Oxford: Blackwell.

Carnap, Rudolf. 1928. *The Logical Structure of the World*. Translated by Rolf A. George. Berkeley, CA: University of California Press.

———. 1932. 'On Protocol Sentences'. Translated by Richard Creath and Richard Nollan. *Noûs* 21 (4): 457–70. https://doi.org/10.2307/2215667.

———. 1934. *Logical Syntax of Language*. Translated by Amethe Smeaton. London: Routledge.

———. 1947. *Meaning and Necessity - a Study in Semantics and Modal Logic*. Chicago, IL: University of Chicago Press.

———. 1950. *Logical Foundations of Probability*. Chicago, IL: University of Chicago Press.

———. 1963. 'Intellectual Autobiography'. In *The Philosophy of Rudolf Carnap*, edited by Paul Arthur Schilpp, 1–83. Open Court.

Carus, A. W. 2007. *Carnap and Twentieth-Century Thought: Explication as Enlightenment*. Cambridge, UK: Cambridge University Press.

Chisholm, Roderick M. 1966. *Theory of Knowledge*. 3rd ed. London: Prentice-Hall International.

Cleve, James Van. 1999. *Problems from Kant*. Oxford: Oxford University Press.

Coffa, J. Alberto. 1991. *The Semantic Tradition from Kant to Carnap: To the Vienna Station*. Edited by Linda Wessels. Cambridge, UK: Cambridge University Press.

Creath, Richard. 2009. 'The Gentle Strength of Tolerance: The Logical Syntax of Language and Carnap's Philosophical Programme'. In *Carnap's Logical Syntax of Language*, edited by Pierre Wagner, 203–14. Houndmills, UK: Palgrave-Macmillan.

Davidson, Donald. 1973. 'On the Very Idea of a Conceptual Scheme'. *Proceedings and Addresses of the American Philosophical Association* 47: 5–20. https://doi.org/10.2307/3129898.

———. 1983. 'A Coherence Theory of Truth and Knowledge'. In *Subjective, Intersubjective, Objective*, 137–53. (2001). Oxford: Clarendon Press.

———. 2001. *Subjective, Intersubjective, Objective*. Oxford: Clarendon Press.

De Pierris, Graciela, and Michael Friedman. 2013. 'Kant and Hume on Causality'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2013. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/win2013/entries/kant-hume-causality/.

DeRose, Keith. 1992. 'Contextualism and Knowledge Attributions'. *Philosophy and Phenomenological Research* 52 (4): 913–29. https://doi.org/10.2307/2107917.

Descartes, René. 1641. *Meditations, Objections, and Replies*. Edited and translated by Roger Ariew and Donald Cress. (2006). Indianapolis, IN: Hackett Publishing Company. https://www.hackettpublishing.com/meditations-objections-and-replies.

Dicker, Georges. 1998. *Hume's Epistemology and Metaphysics: An Introduction*. Abingdon, Oxon: Routledge.

DiSalle, Robert. 2016. 'Space and Time: Inertial Frames'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2016. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/win2016/entries/spacetime-iframes/.

Dreyfus, Hubert L. 2013. 'The Myth of the Pervasiveness of the Mental'. In *Mind, Reason, and Being-in-the-World: The McDowell-Dreyfus Debate*, edited by J. K. Schear, 15–40. Abingdon: Routledge.

Dummett, Michael. 1991. *Frege: Philosophy of Mathematics*. London: Duckworth.

Einstein, Albert. 1916. *Relativity: The Special and the General Theory*. Translated by Robert W. Lawson. (2005). New York: Pi Press.

Evans, Gareth. 1979. 'Reference and Contingency'. *The Monist* 62 (2): 161–89. https://doi.org/10.5840/monist197962220.

Fitch, Gregory W. 2008. 'Analyticity and Necessity in Leibniz'. *Journal of the History of Philosophy* 17 (1): 29–42. https://doi.org/10.1353/hph.2008.0730.

Fodor, Jerry. 1998. *Concepts: Where Cognitive Science Went Wrong*. Oxford: Clarendon Press.

Fodor, Jerry, and Ernie Lepore. 2001. 'Brandom's Burdens: Compositionality and Inferentialism'. *Philosophy and Phenomenological Research* 63 (2): 465–81.

Frege, Gottlob. 1879. 'Begriffsschrift. A Formula Language of Pure Thought Modelled on That of Arithmetic'. In *The Frege Reader*, edited by Michael Beaney. Blackwell.

———. 1884. *The Foundations of Arithmetic*. Translated by J. L. Austin. (1960). New York: Harper & Brothers.

———. 1892. 'Sense and Reference'. *The Philosophical Review* 57 (3): 209–30. https://doi.org/10.2307/2181485.

———. 1893. *Basic Laws of Arithmetic*. Edited and translated by Philip A. Ebert and Marcus Rossberg. Vol. 1. 2 vols. (2013). Oxford: Oxford University Press.

———. 1914. 'Logic in Mathematics'. In *Posthumous Writings*, edited by Hans Hermes, Friedrich Kambartel, and Friedrich Kaulbach, translated by Peter Long and Roger White, 203–50. (1979).

Friedman, Michael. 1987. 'Carnap's Aufbau Reconsidered'. *Noûs* 21 (4): 521–45. https://doi.org/10.2307/2215671.

———. 1999. *Reconsidering Logical Positivism*. Cambridge, UK: Cambridge University Press. https://doi.org/10.1017/CBO9781139173193.

Glock, Hans-Johann. 2003. 'The Linguistic Doctrine Revisited'. *Grazer Philosophische Studien* 66 (1): 141–70. https://doi.org/10.1163/18756735-90000815.

Glüer, Kathrin. 2003. 'Analyticity and Implicit Definition'. *Grazer Philosophische Studien* 66 (1): 37–60.

Goodman, Nelson. 1954. *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press.

———. 1978. *Ways of Worldmaking*. Indianapolis, IN: Hackett Publishing.

Gotterbarn, Donald. 1974. 'Kant, Hume and Analyticity'. *Kant-Studien* 65 (1–4): 274–283. https://doi.org/10.1515/kant.1974.65.1-4.274.

Grice, H. P., and P. F. Strawson. 1956. 'In Defense of a Dogma'. *The Philosophical Review* 65 (2): 141–58. https://doi.org/10.2307/2182828.

Hacker, P. M. S. 1986. *Insight and Illusion: Themes in the Philosophy of Wittgenstein*. 2nd ed. Oxford: Clarendon Press.

———. 1996. *Wittgenstein's Place in Twentieth-Century Analytic Philosophy*. Oxford: Blackwell.

———. 2009. 'A Philosopher of Philosophy'. *The Philosophical Quarterly* 59 (235): 337–48.

Hale, Bob, and Crispin Wright. 2000. 'Implicit Definition and the a Priori'. In *New Essays on the a Priori*, edited by Paul Artin Boghossian and Christopher Peacocke, 286–319. Oxford: Clarendon Press.

Hanfling, Oswald. 1981. 'Leibniz's Principle of Reason'. In *Gottfried Wilhelm Leibniz: Metaphysics and Its Foundations I*, edited by R. S. Woolhouse, 1:74–81. (1994). London: Routledge.

———. 2000. *Philosophy and Ordinary Language: The Bent and Genius of Our Tongue*. London: Routledge.

Hume, David. 1740. *A Treatise of Human Nature*. Edited by L. A. Selby-Bigge and P. H. Nidditch. Oxford: Oxford University Press.

———. 1777. *Enquiries Concerning the Human Understanding and Concerning the Principles of Morals*. Edited by L. A. Selby-Bigge and P. H. Nidditch. Oxford: Clarendon Press.

Hylton, Peter. 2016. 'Willard van Orman Quine'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2016. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/win2016/entries/quine/.

Jackson, Frank. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.

James, William D. 1890. *The Principles of Psychology*. Cambridge, MA: Harvard University Press.

Juhl, Cory, and Eric Loomis. 2009. *Analyticity*. Abingdon, Oxon: Routledge.

Kanas, Nick, and Dietrich Manzey. 2008. *Space Psychology and Psychiatry*. The Space Technology Library 22. Dordrecht: Springer Netherlands. https://doi.org/10.1007/978-1-4020-6770-9_2.

Kant, Immanuel. 1763. 'Attempt to Introduce the Concept of Negative Magnitudes into Philosophy'. In *Theoretical Philosophy, 1755-1770*, translated by David Walford. Cambridge, UK: Cambridge University Press.

———. 1783. *Prolegomena to Any Future Metaphysics: And the Letter to Marcus Herz, February 1772*. Indianapolis: Hackett.

———. 1787. *Critique of Pure Reason, Second Edition*. Translated by Norman Kemp Smith. Hampshire, UK: Palgrave Macmillan.

———. 1880. 'The Jäsche Logic'. In *Lectures on Logic*, translated by J. Michael Young, 356–87. Cambridge, UK: Cambridge University Press.

Kanterian, Edward. 2012. *Frege: A Guide for the Perplexed*. London: Continuum.

Kaplan, David. 1979. 'On the Logic of Demonstratives'. *Journal of Philosophical Logic* 8 (1): 81–98. https://doi.org/10.1007/BF00258420.

Knobe, Joshua, and Shaun Nichols, eds. 2007. *An Experimental Philosophy Manifesto*. Oxford: Oxford University Press.

Kornblith, Hilary. 2002. *Knowledge and Its Place in Nature*. Oxford: Clarendon.

Körner, Stephan. 1955. *Kant*. Harmondsworth, UK: Penguin Books.

Kremer, Michael. 2010. 'Representation or Inference: Must We Choose? Should We?' In *Reading Brandom: On Making It Explicit*, edited by Bernhard Weiss and Jeremy Wanderer, 227–46. Abingdon: Routledge.

———. 2016. 'A Capacity to Get Things Right: Gilbert Ryle on Knowledge'. *European Journal of Philosophy*, January, 1–22. https://doi.org/10.1111/ejop.12150.

Kripke, Saul. 1972. *Naming and Necessity*. Oxford: Blackwell.

Kroedel, Thomas. 2012. 'Implicit Definition and the Application of Logic'. *Philosophical Studies* 158 (1): 131–48. https://doi.org/10.1007/s11098-010-9675-0.

Künne, Wolfgang. 2006. 'Analyticity and Logical Truth: From Bolzano to Quine'. In *The Austrian Contribution to Analytic Philosophy*, edited by Mark Textor, 184–249. Abingdon, Oxon: Routledge.

Lance, Mark Norris, and John O'Leary-Hawthorne. 1997. *The Grammar of Meaning: Normativity and Semantic Discourse*. Cambridge, UK: Cambridge University Press.

Leibniz, Gottfried Wilhelm Freiherr von. 1686. 'Discourse on Metaphysics'. In *Philosophical Texts*, translated by Richard Francks and R. S. Woolhouse. (1998). Oxford: Oxford University Press.

———. 1714. 'Monadology'. In *Philosophical Texts*, translated by Richard Francks and R. S. Woolhouse, 267–81. (1998). Oxford: Oxford University Press.

———. 1973. *Philosophical Writings*. Edited and translated by G. H. R. Parkinson and Mary Morris. London: Dent.

Levin, Janet. 2009. 'Experimental Philosophy'. *Analysis* 69 (4): 761–69. https://doi.org/10.1093/analys/anp104.

Locke, John. 1689. *An Essay Concerning Human Understanding*. Edited by Kenneth Winkler. 1996. Indianapolis, IN: Hackett Publishing Company.

Maher, Chauncey. 2012. *The Pittsburgh School of Philosophy: Sellars, McDowell, Brandom*. London: Routledge.

Mates, Benson. 1986. *The Philosophy of Leibniz: Metaphysics and Language*. Oxford: Oxford University Press.

McDowell, John. 1994. *Mind and World: With a New Introduction*. Cambridge, MA: Harvard University Press.

———. 1995. 'Reply to Gibson, Bryne, and Brandom'. *Philosophical Issues* 7: 283–300.

———. 1997. 'Brandom on Representation and Inference'. *Philosophy and Phenomenological Research* 57 (1): 157–62.

———. 1998. 'Reply to Commentators'. *Philosophy and Phenomenological Research* 58 (2): 403–31.

———. 2005. 'Motivating Inferentialism: Comments On'. *Pragmatics & Cognition* 13 (1): 121–40.

———. 2009. 'Conceptual Capacities in Perception'. In *Having the World in View: Essays on Kant, Hegel, and Sellars*, 127–44. Cambridge, MA: Harvard University Press.

———. 2010. 'Brandom on Observation'. In *Reading Brandom: On Making It Explicit*, edited by Bernhard Weiss and Jeremy Wanderer, 129–44. Abingdon: Routledge.

———. 2013. 'The Myth of the Mind as Detached'. In *Mind, Reason, and Being-in-the-World: The McDowell-Dreyfus Debate*, edited by J. K. Schear, 41–58. Abingdon: Routledge.

Meeker, Kevin. 2007. 'Hume on Knowledge, Certainty and Probability: Anticipating the Disintegration of the Analytic/Synthetic Divide?' *Pacific Philosophical Quarterly* 88 (2): 226–242.

Moore, A. W. 2009. 'Not to Be Taken at Face Value'. *Analysis* 69 (1): 116–25. https://doi.org/10.1093/analys/ann040.

Murzi, Julien, and Florian Steinberger. 2016. 'Inferentialism'. In *A Companion to the Philosophy of Language*, edited by Bob Hale, Alex Miller, and Crispin Wright, 2nd ed. London: Blackwell.

NASA. 2015. 'What Is Microgravity?' NASA. 13 May 2015. http://www.nasa.gov/centers/glenn/shuttlestation/station/microgex.html.

OED. 2017. 'Bachelor, N.' *OED Online*. Oxford University Press. http://www.oed.com/view/Entry/14313.

Owen, David. 1999. *Hume's Reason*. Oxford: Oxford University Press.

Papineau, David. 2015. 'The Nature of a Priori Intuitions: Analytic or Synthetic?' In *Experimental Philosophy, Rationalism, and Naturalism: Rethinking Philosophical Method*, edited by Eugen Fischer and John Collins, 51–71. Abingdon, Oxon: Routledge.

Pereboom, Derk. 2014. 'Kant's Transcendental Arguments'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2014. US: Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/fall2014/entries/kant-transcendental/.

'Peter Hacker Did Not like Timothy Williamson's Book'. 2009. Leiter Reports: A Philosophy Blog. 4 May 2009. https://web.archive.org/web/20090504001652/http://leiterreports.typepad.com/blog/2009/04/peter-hacker-did-not-like-timothy-williamsons-book-1.html.

Pössel, Markus. 2005. 'The Dialectic of Relativity'. *Einstein Online* 01 (1005). http://www.einstein-online.info/spotlights/dialectic.

———. 2010. 'Is the Whole the Sum of Its Parts?' *Einstein Online* 04 (1003). http://www.einstein-online.info/spotlights/binding_energy.

Putnam, Hilary. 1962. 'The Analytic and the Synthetic'. In *Philosophical Papers: Volume 2: Mind, Language and Reality*, 2:33–69. (1975). Cambridge: Cambridge University Press.

———. 1975. 'The Meaning of "Meaning"'. In *Philosophical Papers: Volume 2: Mind, Language and Reality*, 2:215–71. Cambridge: Cambridge University Press.

Quine, Willard Van Orman. 1951. 'Two Dogmas of Empiricism'. In *From a Logical Point of View: 9 Logico-Philosophical Essays*, 20–46. 1961. Cambridge, MA: Harvard University Press.

———. 1962. 'Carnap and Logical Truth'. *Synthese* 12 (4): 350–74. https://doi.org/10.1007/BF00485423.

Raffman, Diana. 2005. 'How to Understand Contextualism about Vagueness: Reply to Stanley'. *Analysis* 65 (287): 244–48. https://doi.org/10.1111/j.1467-8284.2005.00558.x.

Redding, Paul. 2007. *Analytic Philosophy and the Return of Hegelian Thought*. Cambridge, UK: Cambridge University Press.

Reinach, Adolf. 1908. 'Kant's Interpretation of Hume's Problem'. Translated by J. N. Mohanty. *The Southwestern Journal of Philosophy* 7 (2): 161–88.

Rescher, Nicholas. 2011. 'Leibniz On Infinite Analyticity'. In *Philosophical Episodes*, 27–31. Frankfurt: Ontos-Verl. https://doi.org/10.1515/9783110319934.27.

Richardson, Alan W. 1998. *Carnap's Construction of the World: The Aufbau and the Emergence of Logical Empiricism*. Cambridge, UK: Cambridge University Press.

Rorty, Richard, ed. 1967. *The Linguistic Turn: Essays in Philosophical Method*. Chicago: University of Chicago Press.

Rundle, Bede. 1979. *Grammar in Philosophy*. Oxford: Clarendon Press.

Russell, Bertrand. 1905. 'On Denoting'. *Mind* 14 (56): 479–93.

———. 1956. 'The Philosophy of Logical Atomism'. In *Logic and Knowledge: Essays 1901-1950*, edited by Robert Charles Marsh, 175–284. London: Routledge.

Russell, Gillian. 2007. *Truth in Virtue of Meaning: A Defence of the Analytic/Synthetic Distinction*. Oxford: Oxford University Press.

———. 2010. 'A Review of Timothy Williamson's the Philosophy of Philosophy'. *Philosophical Books* 51 (1): 39–52. https://doi.org/10.1111/j.1468-0149.2010.00500.x.

Ryle, Gilbert. 1933. 'About'. In *Collected Papers Vol 2: Collected Essays, 1929-1968*, 86–88. Abingdon: Routledge.

———. 1945. 'Philosophical Arguments'. In *Collected Papers Vol 2: Collected Essays, 1929-1968*, 203–21. Abingdon: Routledge.

———. 1949. *The Concept of Mind: 60th Anniversary Edition*. (2009). London: Routledge.

———. 1962. 'Abstractions'. In *Collected Papers Vol 2: Collected Essays, 1929-1968*, 448–58. 2009. London: Routledge.

Sandis, Constantine. 2010. 'The Experimental Turn and Ordinary Language'. *Essays in Philosophy* 11 (2).

———. 2011. 'Review of Gilbert Ryle, The Concept of Mind, Collected Papers Volume I, Collected Papers Volume II'. *Philosophy in Review* 31 (6): 455–57.

Schlick, Moritz. 1925. *General Theory of Knowledge*. Translated by Albert E. Blumberg. 2nd ed. New York: Springer-Verlag.

Schroeder, Severin. 2006. *Wittgenstein: The Way out of the Fly-Bottle*. Cambridge, UK: Polity Press.

———. 2009. 'Analytic Truths and Grammatical Propositions'. In *Wittgenstein and Analytic Philosophy: Essays for P.M.S. Hacker*, edited by P. M. S. Hacker, Hans-Johann Glock, and John Hyman, 83–108. Oxford: Oxford University Press.

Sellars, Wilfrid. 1953. 'Inference and Meaning'. *Mind* 62 (247): 313–38.

———. 1956. 'Empiricism and the Philosophy of Mind'. In *Knowledge, Mind, and the Given*, 205–76. Indianapolis, IN: Hackett Pub.

Spinoza, Benedictus de. 1663. 'Parts I and II of Descartes' Principles of Philosophy'. In *The Collected Works of Spinoza*, edited and translated by Edwin Curley. Vol. 1. Princeton, NJ: Princeton University Press.

Stanley, Jason. 2003. 'Context, Interest Relativity and the Sorites'. *Analysis* 63 (4): 269–80. https://doi.org/10.1093/analys/63.4.269.

Strawson, P. F. 1950. 'On Referring'. *Mind* 59 (235): 320–44.

Suchting, W. A. 1966. 'Hume and Necessary Truth'. *Dialogue: Canadian Philosophical Review / Revue Canadienne de Philosophie* 5 (1): 47–60. https://doi.org/10.1017/S0012217300036210.

Tanney, Julia. 2013. *Rules, Reason, and Self-Knowledge*. Cambridge, MA: Harvard University Press.

Thomasson, Amie. 2007. *Ordinary Objects*. Oxford: Oxford University Press.

Travis, Charles. 2008a. 'Introduction'. In *Occasion-Sensitivity: Selected Essays*. Oxford: Oxford University Press.

———. 2008b. *Occasion-Sensitivity: Selected Essays*. Oxford: Oxford University Press.

Weinberg, Jonathan M., Shaun Nichols, and Stephen Stich. 2001. 'Normativity and Epistemic Intuitions'. *Philosophical Topics* 29 (1/2): 429–60. https://doi.org/10.5840/philtopics2001291/217.

Weiner, Joan. 2007. 'What's in a Numeral? Frege's Answer'. *Mind* 116 (463): 677–716. https://doi.org/10.1093/mind/fzm677.

Whitehead, Alfred North, and Bertrand Russell. 1910. *Principia Mathematica*. Cambridge, UK: Cambridge University Press.

Whiting, Daniel. 2007. 'The Normativity of Meaning Defended'. *Analysis* 67 (294): 133–40. https://doi.org/10.1111/j.1467-8284.2007.00663.x.

Wiggins, David. 2001. *Sameness and Substance Renewed*. Cambridge, UK: Cambridge University Press.

Williamson, Timothy. 1994. *Vagueness*. London: Routledge.

———. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.

———. 2007. *The Philosophy of Philosophy*. Oxford: Blackwell.

———. 2012. 'Boghossian and Casalegno on Understanding and Inference'. *Dialectica* 66 (2): 237–47. https://doi.org/10.1111/j.1746-8361.2012.01295.x.

Wilson, Mark. 1988. 'Nature's Demands on Language'. *Philosophical Topics* 16 (1): 285–336. https://doi.org/10.5840/philtopics198816110.

———. 2006. *Wandering Significance: An Essay on Conceptual Behavior*. Oxford: Clarendon Press.

Wittgenstein, Ludwig. 1921. *Tractatus Logico-Philosophicus*. London: Routledge.

———. 1953. *Philosophical Investigations*. Translated by G. E. M. Anscombe, P. M. S. Hacker, and Joachim Schulte. 4th ed. Chichester: Wiley-Blackwell.

———. 1956. *Remarks on the Foundations of Mathematics*. Edited by Georg Henrik von Wright, Rush Rhees, and G. E. M. Anscombe. Translated by G. E. M. Anscombe. 2nd ed. (1967). Cambridge, MA: The M.I.T. Press.

———. 1958. *The Blue and Brown Books*. 2nd ed. (1969). Oxford: Blackwell.

———. 1974. *Philosophical Grammar*. Edited by Rush Rhees. Translated by Anthony Kenny. Oxford: Basil Blackwell.

Woolhouse, R. S. 1998. 'Editor's Introduction'. In *Philosophical Texts*, 5–49. Oxford: Oxford University Press.

Wright, Georg Henrik von. 1963. *Norm and Action: A Logical Enquiry*. London: Routledge & Kegan Paul.