University of Reading

Sailing Through Narrow Straits: Necessity, Contingency, and Language

PhD in Philosophy

Department of Philosophy

Sam W. A. Couldrick

September 2023

Sailing Through Narrow Straits: Necessity, Contingency, and Language

Sam W. A. Couldrick

Abstract

This thesis examines necessary truth and defends a normative, or linguistic, account of it. Roughly, it holds that necessary truths state or follow from conceptual norms (i.e., norms that determine patterns of correct concept use). While the thesis touches upon logical and mathematical truth, its primary focus are those necessary truths typically expressed using natural language. The thesis has three parts.

In Part I, I criticise metaphysical accounts of necessity and present and defend a normative account of it. At no point do I give a history of normative accounts, but clearly their roots are to be found in the first half of the twentieth century – in the works of Wittgenstein and Carnap, for example.

In Part II, I consider whether language can sustain the normative account. Some argue that the account requires language to be regimented in a way that it is not. I show that while it requires a distinction in kind between empirical and conceptual principles, it nevertheless makes room for indeterminacy regarding whether a given statement is an empirical claim or follows from conceptual norms.

Finally, in Part III, I consider the relationship between the world and our conceptual scheme. I argue that denying our concepts answer to the world does not mean that they cannot be justified. The normative account does not say that we have no reasons for categorising things in a certain way, but rather that natural facts, in combination with our interests, are fit to provide them.

The purpose of the thesis is to show that normative accounts of necessity can be much more robust than they are often given credit for and needn't have the malign implications often associated with them.

Preface

Declaration: I confirm that this is my own work and the use of all material from other sources has been properly and fully acknowledged.

Sam Couldrick

Acknowledgements: I first began to think about the topics in this thesis during my undergraduate degree at Oxford and I would like to thank many of the academics there that supported, stimulated, and challenged my thinking in various ways. These include Dave Leal, Brian King, Chris Timpson, and Ralph Walker.

I returned to these topics throughout the course of the BPhil where again Ralph's supervisions were of immense value to me. I would also like to thank my BPhil thesis supervisor, Adrian Moore, who was incredibly supportive. It was in that thesis that I first played around (somewhat naively) with some of the ideas herein, and Adrian's insights during that time have been useful here too.

Of course, most important to this project has been my supervisor, Severin Schroeder, who I would like to thank both for his encouragement throughout the writing of this thesis and the countless insightful comments he has made on its contents. I am also grateful to Max de Gaynesford for his feedback on a near-final copy of the thesis and to those across the faculty and student body that I have discussed these topics with.

Finally, I would like to thank my family and in particular my parents for their unwavering support throughout this process. I know that they are as relieved as I am that it is coming to a close.

Contents

Introduction

Necessary truths are said to be true in all circumstances, regardless of context. Typical examples are '2 + 2 = 4', 'bachelors are unmarried men', and 'something cannot be red and green all over simultaneously'. Necessary truths exhibit a unique kind of generality and have long been subject to the fascination of philosophers. They appear to be distinct from empirical generalities, such as those discovered in science, for their refutation is impossible. It is generally accepted that empirical truths could have turned out to be false. Not so for necessary truths.

Our interest lies in what we might call the 'absolute', or 'basic', case of necessity. After all, there are many circumstances in which we say something must be the case when, in the strictest sense, it doesn't have to be. For example, it may be true that something must happen, *given certain laws of nature*. Were the laws of nature different, however, what is now physically impossible may not be. Likewise, someone might say something 'can't be true', *given what they know*. But what is thereby ruled out could nevertheless have been the case (they just know it not to be). Similarly, it might be said that I cannot take a particular course of action, *while following a certain moral framework (i.e., because it would be wrong for me to do so)*. Yet I don't have to do what I ought to. In each of these cases, the italicised text needn't be made explicit, but it underwrites the relevant statements. In that sense, we might think of these necessities as conditional. Once we include their conditions, we arrive at the basic case.

Necessary truths play a central role in philosophy. Philosophers are often interested in the nature or essence of things. While scientists might attempt to discover laws of the world we inhabit, philosophers are sometimes said to consider all the worlds we might have inhabited. As Bertrand Russell put it, 'a philosophical proposition must be applicable to everything that exists or may exist.'[1] Irrespective of whether that is a fair reflection of all 'philosophical propositions', this level of generality is somewhat characteristic of philosophy. There is, however, significant disagreement about the nature of necessity and thus the nature of philosophy itself.

On what has become an orthodox understanding of philosophy (at least within Western philosophical quarters), mapping out (some) necessary truths is a matter of tracing the contours of reality. Some necessary propositions describe the world. The world, in at least some instances, makes certain necessary propositions true. For this form of 'realism', there are some necessary propositions that are analogous to empirical ones, only the features of the world that make

[1] Russell (1914 [1917], 110).

necessary propositions true are in some way firmer or more general than those that make empirical propositions true. Some philosophy, on this picture, is a matter of trying to describe the 'fundamental nature of reality' or discover 'real essences' – a physics for a level of reality distinct from that studied by physics: a metaphysics.

But, as Kant explained, there is an immediate problem with this. On the one hand, metaphysics purports to establish knowledge independent of experience. On the other, it is said to yield knowledge of reality. That is, it lays claim to knowledge of the world without grounding it on the mode of access we use to learn about reality, namely experience. One interpretation of Kant's solution to this problem is that, rather than tracing the contours of a mind-independent reality, necessary truths instead explicate certain 'transcendental' mental structures. But it isn't clear to what extent this can help. How do we know such structures exist and is it obvious we would have access to them? Are they not, in the end, still constituents of reality and is our *a priori* access to them not still opaque?

Christopher Peacocke calls this the 'integration challenge'.[2] For an account of these truths to be plausible, we need to be able to reconcile what is involved in the truth of the relevant statements with a credible epistemological explanation of how we come to know them. One might say we must reconcile the metaphysics of necessity with its epistemology. Yet it is far from clear that the orthodox account of necessity can provide a credible epistemology. How to determine, for example, what properties something *must* have? As Hume pointed out, this is not to be learnt from experience: that something has always been the case does not mean that it must be. It is far from clear that we have a reliable faculty for establishing metaphysically necessary truths.[3]

One reaction to the integration challenge is to expunge the notion of necessary truth altogether. Yet we do seem capable of distinguishing between relevant cases. Compare, for example, a basic mathematical truth with the fact that the sun has always risen. Moreover, discarding necessity because it cannot be empirically verified seems misguided. In short, before suggesting that we can't get to the necessary from the empirical – as though it is a realm beyond our reach – we should interrogate the role modals have in our empirical discourse. For the discussion of conditional necessities above suggests basic necessity precedes (perhaps, makes possible) the very empirical claims we are supposedly restricted to. They do so by articulating what is involved in a given claim. For example, by acknowledging that it can't have been Ayla at

[2] Peacocke (1999, 1).
[3] Cf. Nozick (2001, 120-2) and Thomasson (2021).

the park because she is in Spain, I thereby articulate (part of) what is involved in the claim that she was at the park and the claim that she was in Spain. Their being incompatible partly determines the claim being made and the places being referred to. This is at least a *prima facie* reason to resist the crudest of empiricist urges to deny necessity altogether.[4]

The metaphysical account of necessity correctly eschews modal scepticism but nevertheless struggles to answer the integration challenge. The epistemology appears to rest on a form of 'rational insight', or, in contemporary philosophical parlance, 'intuitions' (intellectual seemings). But these often lack a credible foundation. The point here is not to cast doubt on the very notion of an intuition. There is an ordinary sense in which we know ourselves to have intuitions – that is, instincts on a subject matter that are not the result of conscious reasoning or recall. However, there is usually a basis for such intuitions, or at least there needs to be for them to be remotely credible. Someone can make an empirical claim, for example, that might immediately sound suspect to those familiar with the area of research. While this intuition is, for the time being, not supported by evidence, it may well be founded on the researcher's familiarity with the subject. But a story like this cannot be told in the case we are interested in. The point is that our experience is not fit to settle the questions being asked, and so the grounds that support our intuitions in other cases cannot support them here. Similarly, it's one thing to use counterfactuals to describe how the world works, supported by inductive evidence, observations, and verifiable claims. Quite another to suppose that same faculty could obtain knowledge about which aspects of those workings are *necessary*. The fundamental question that still needs to be answered is this: what connects our thinking, or intuitions, to a level of reality we cannot investigate via experience? This is altogether mysterious and naturalistic explanations are not forthcoming.[5]

Proponents of metaphysics are no strangers to this quandary. Theodore Sider, for example, suggests the epistemology is 'far from clear; this any metaphysician should concede.'[6] Likewise, Timothy Williamson, having castigated 'appeals to the authority of Kant' as ringing hollow, admits that 'we do not fully understand how thinking can provide new knowledge', but adds that 'the cases of logic and mathematics constitute overwhelming evidence that it does so'.[7] This gentle rejoinder, however, obscures what is really at issue, as Hanjo Glock explains.

[4] Another way of putting this is to first recognise the role necessity plays in structuring counterfactual thinking and second recognise the role counterfactuals play in everyday discourse. Cf. Godman, Mallozzi and Papineau (2020) and Williamson (2007).
[5] Cf. Warren (2022, 27-8).
[6] Sider (2011, 12).
[7] Williamson (2004, 111; 127).

> While logic and mathematics demonstrate the possibility of non-trivial *a priori* truths—a point Kant himself emphasized through his notion of the synthetic *a priori*—it is far from obvious, to put it mildly, that they provide *a priori* knowledge of reality. (Glock 2012a, 409)

The integration challenge may not be fatal for metaphysical necessity. It could be that we have ruled out modal scepticism and that alternative accounts of necessity, which perhaps meet the integration challenge better, are unsatisfactory for different reasons. But I will argue that the metaphysical account is not the best available and an alternative is more compelling.

In Part I, I criticise metaphysical accounts of necessity and discuss and defend one alternative that others have, in exactly the fashion just described, discarded. The normative (or linguistic) account of necessity treats necessary truths as having a conceptual (or semantic) foundation. The pursuit of necessary truths is not a matter of tracing the contours of reality, but rather constitutes an investigation into our means of representing reality. Whereas some think Gettier's thought experiments reveal truths about not just our concept 'knowledge' but *knowledge* itself, a linguistic account considers the distinction to be chimerical.[8] Knowledge just is that which falls under our concept, and so an investigation into the concept 'knowledge' is an investigation into the nature of *knowledge*.

Having explained how conceptual norms might account for necessity, in Part II I describe the role necessity plays in language. One reason for doing this is to counter the suggestion that the normative account of necessity requires that language be regimented in a way that it is not. I show that the normative account doesn't imply that language is constituted by a strict set of rules, and therefore doesn't require that we can, in all circumstances, determine whether a given statement is empirical or conceptual in nature.

Finally, in Part III, I describe the relationship between the world and our conceptual scheme and the necessary truths that help constitute it. Denying that our concepts answer to the world does not mean that they cannot be said to be justified when they meet our interests. To meet our interests, they must pay attention to how the world is. In other words, the normative account does not say we have no reasons for categorising things in a certain way, but rather that we can appeal to empirical facts to explain this. It doesn't rely on our concepts mirroring some metaphysical order. The realm of natural facts is quite sufficient.

[8] Gettier (1963).

The course this account charts is narrow.[9] On the one hand, it claims that necessary truths express conceptual norms that we can choose to accept or reject. On the other, it suggests that this seemingly contingent basis does not jeopardise necessity or amount to modal scepticism. It holds that there is fundamentally a distinction in kind between empirical and conceptual propositions, but that it may not be determinate what the nature of a given proposition is. And it suggests that while conceptual propositions do not answer to the world, they may nevertheless be justified – in part – by that world.

The purpose of explaining its subtleties and showing how it can answer some leading objections is to suggest that the account is more attractive than is often acknowledged – that renewed interest in it is warranted. In this way, the thesis connects to, and builds on, some recent work that has sought to revive normative accounts of necessity by casting doubt on the force of objections often made against it.[10] The thesis also draws a great deal from the philosophical work of historical proponents of similar accounts, as well as more recent scholarship. The conclusion is that normative accounts of necessity can be much more robust than they are often given credit for and needn't have the malign implications often associated with them.

[9] The title of this thesis is inspired by a passage in Dilman (2002, 57-8).
[10] Of particular note is Thomasson (2020a).

Part I

In what follows, I consider the case for necessary truths being substantial and argue that it is unconvincing. I suggest instead that a semantic understanding of these truths is more compelling. A sketch of the account is provided, and objections raised. Throughout there are examples of necessary truths, though none are particularly central to my argument. The question is how to account for necessary truth, not what the necessary truths are. For the time being, it is assumed that necessary truths exist. Any concern that I am insensitive to the possibility that the semantics involved are more complex than I suggest should be assuaged by Part II.

(1) Necessity and Conceptual Limits

One way to think about necessary truths is that they are limit-setting in that they mark the limits of what is possible. Any proposition that denies a necessary truth is necessarily false, impossible. A question that has exercised philosophers for centuries, and that was brought to a head in Kant's magnum opus, *Critique of Pure Reason*, is whether these limits are set by reality, us, or something else. If set by reality, then there is a sense in which necessary truths are *about* or *describe* the world. If I claim a proposition to be necessary, I thereby lay down a condition that I believe reality *must* meet, and it is a necessary *truth* just in case the world must, in fact, meet that condition. Thus, a necessary truth could be said to impose a genuine constraint on reality.

This chapter is concerned with precisely that notion. Can the thought that these limits are genuine constraints on reality survive close inspection? In the first section, I consider what I take to be an indisputable feature of necessary propositions, namely that they set limits on the concepts contained within them. In the second part, I argue that necessary propositions cannot be thought of as imposing genuine constraints on reality and that the role they have in setting conceptual limits goes some way to explaining why. In the third and fourth sections, I consider some objections. Throughout the chapter, an alternative account of necessary truth is suggested. Only a rough picture is offered, but it is developed and defended throughout the rest of the thesis.

1. Conceptual limits

The nature of necessity, and its explanation, can begin to be elucidated by considering what necessary truths serve to do. Take the following example:

i) A square is a shape with sides of equal length.

For this to be a necessary truth, what must hold of any given square? What this unambiguously tells us is that each of its sides must be of equal length. That is, whatever falls under the concept 'square' must have this feature. It – of course – does not say that everything with this feature is a square. In other words, this necessary truth states a necessary (but not sufficient) condition of squarehood. This condition sets a limit on what does and does not count as a square. To know the above is to know that, whatever shape a square is, its sides must be equal in length. We can imagine someone whose knowledge of squares stretches only as far as the above condition, ruling certain shapes out whilst keeping others in contention. For instance, were they to

encounter a scalene triangle, they would know that whatever that shape is, it isn't a square. Already, then, we can begin to see what necessary truths serve to do: they limit the use of the concepts in question. In so doing, they delimit the boundaries of those concepts, determining what counts as being an instantiation of them. They serve to differentiate the concepts in question from others.

In so differentiating the concepts, they contribute to an explanation of the meaning of the given terms. As we know, squares are shapes with four sides of equal length. This is (part of) what 'square' means. The truths that are said to be necessary of the square are partly constitutive of its meaning. Without them, 'square' would mean something else – or nothing at all. Thus, necessary truths, which state conditions under which it is correct to apply the relevant concepts, provide characteristics we use to categorise, classify, and describe the world. These truths help to isolate objects in, and features of, the world we experience, enabling us to represent that world to ourselves and each other (and much else besides, including expressing joy, asking questions, marrying our beloved, etc.).

The conceptual contribution these truths make is indispensable.[11] I cannot make any headway in thinking about the world without some concepts with which to do it. Such concepts, if they are to be at all useful, must distinguish one thing from the rest. By this I do not mean that we have no use for the concept 'everything', but rather that its use comes from its contrast with something else. The point here is merely that if there is no sense in the negation of some concept, then that concept has so far failed to differentiate its extension from anything else – and is, in this way, uninformative when used.[12] In applying such a concept, we would fail to bring anything in particular to attention. Thus, anyone wishing to, for example, propose an account of our relationship with objects in the world as being fundamentally indirect in nature must – for us to make anything of such a claim – provide us with an idea of what direct access to the world might look like (what it is to be with and without such access). Without this contrast, the original

[11] Note that I say the conceptual contribution is indispensable, not that simple, necessary conditions are indispensable to concepts. Certain concepts may admit of no such conditions. Moreover, rules admitting of exceptions can still make conceptual contributions. See Part II.

[12] This is not to sign up to the erroneous argument made by Davidson (1973, 20) in relation to the conceptual scheme. Davidson suggests the notion of a conceptual scheme is unintelligible because there cannot be a plurality of such schemes. But this is to assume that for a concept to be meaningful it must be distinguishable from something of the same kind, in this case, another conceptual scheme. But in truth it only needs to be distinguishable from something else. If there could not be more than one universe, 'universe' would not cease to be useful. It would, for example, distinguish itself from 'galaxy', 'planet', 'bookshop'. See Glock (2009, 658).

claim is not in the least bit informative.[13] 'Indirect access' simply stands in the place of 'access', without any material difference between the two (save for the extra word).[14]

The point here is not intended to be limited only to those necessary truths we assent to. Rather, any statement characterised by necessity does itself begin to carve the contours of the concepts contained within it. For instance, imagine someone said the following:

ii) Every man is subject to fits of rage.

If this statement is supposed to exhibit necessity in the relevant sense, then a counterexample will not be countenanced.[15] The reason for this has nothing to do with the psychological hold apparently exceptionless rules have over us, but rather that, quite simply, one cannot be a 'man' (according to this conception) unless one is subject to fits of rage. To count as a 'man', something must be subject to such rage. To try and provide a counterexample would be to miss the point. This statement has a stipulative effect of the following kind: 'if you want to speak as I do, then you must only apply "man" to those subject to fits of rage'. Of course, this is a quite different understanding of 'man' from ours, but the statement is nonetheless true of this different concept. The statement provides an incomplete picture of a concept that could be used to think about and describe the world. If they were claiming that this captures the meaning of our term, then counterexamples (with regards to how we use 'man') would of course be countenanced.

The statements in question, then, have two sides to them. On the one hand, they have a descriptive effect: a square *is* a shape of four sides with equal length. This *is* (part of) what we mean by 'square'. On the other hand, they have a stipulative aspect: a square *must* have four sides with equal length. To speak as we do, to mean the same thing as we do by 'square', you must only use this word with respect to shapes that have four sides of equal length. Of course, there is nothing to prevent anyone from using 'square' to mean something else, but they wouldn't be speaking as we do.

The purpose of what I have so far said is to show that when a statement is said to exhibit necessity it cannot but begin to define the boundaries of the concepts contained within it. That may be one effect amongst many, but it is certainly one effect. This is true even when someone claims that something is necessary when we wouldn't ordinarily say so, or when it isn't a necessity we would ordinarily assent to, perhaps because they are using the same words to mean

[13] Mulhall (2018) makes this criticism in his review of Harman (2018).
[14] Cf. PI, 13.
[15] Cf. Bangu (2019, 449).

something slightly (or significantly) different. For now, I make no submission either way as to whether these novel ways of representing things are constituted by necessary *truths*, but it is true of our imagined person's concept of 'man' that men must be subject to fits of rage. It is, however, false that our ordinary concept of 'man' requires all objects fitting under it to be subject to fits of rage, which is to say that the above statement is not an accurate reflection of how we speak. Frank Jackson makes a similar point with respect to those who think knowledge survives Gettier cases: they may be correct so far as their concept 'knowledge' goes.

It is worth pausing here to clarify something. I have suggested that necessary statements express some limits of the concepts contained within them. But one might agree that it certainly seems the case for some of those concepts, but not for others. Take:

iii) *All bachelors are unmarried men.*

'This is a definition of 'bachelor', but what has it to do with 'man'?', one might ask. Well, the fact that if x is a bachelor, then x is a man, indicates a limit of our concept 'man', as well as that of 'bachelor'. The counting of men in a room, for example, must include all bachelors. Thus, while *iii)* tells us relatively little about the concept 'man', especially compared with that of 'bachelor', it may yet act as a kind of thinking-guide with respect to the concept 'man', which might be useful were we prone to mistakenly excepting bachelors from our counting of men.[16] (Perhaps I am a social scientist who tends to make claims about men despite my only studying the attitudes of married ones.)[17]

The question I am interested in is whether there is a sense of truth in necessary propositions which goes beyond that of the correct codification of an existing (or imagined) language. Can we make sense of the idea that these statements, that express some limits of the concepts contained within them, answer to reality? This picture brings these statements closer to our empirical generalisations, but rather than seeking to describe the physical laws of our world, instead they come to describe some structure that is more fundamental, that would apply to all possible worlds. Notice that, insofar as one accepts the role these statements have in setting limits on the concepts in question, this picture suggests that there is a privileged set of concepts that depicts the world more accurately than the alternatives. Reality would justify the use of some concepts over others, thereby making the statements which describe the meaning of those concepts answer to reality. Those concepts that most accurately depict reality have been said to carve nature at its joints. If this characterisation could be made sense of and were correct, we

[16] See Diamond (2019) for more on thinking-guides.
[17] Büttner (2015) draws a similar distinction between the meanings of terms and definitions.

would have a third aspect to the above kinds of statements. Alongside describing a given linguistic practice and stipulating how one must speak to be aligned with it, necessary propositions might also: *describe some features of reality*.

In *Writing the Book of the World*, Sider presents an account that incorporates something of this view. We see Sider effectively recognise the three aspects of necessary truths that I outlined above in his discussion of analyticity. He mentions, for example, 'the fortuitous convergence of definitionality and joint-carving' by which he means instances where the definition of a concept accurately depicts the structure of the world.[18]

> The world has a distinguished structure, a privileged description. For a representation to be fully successful, truth is not enough; the representation must also use the right concepts, so that its conceptual structure matches reality's structure. There is an objectively correct way to 'write the book of the world'. (Sider 2011, vii)

For our purposes, this structure might be said to ground, or make true, necessary truths. According to such accounts, what matters is whether our network of concepts corresponds with the world's structure. This can be difficult to assess. The world's structure is not tangible. Indeed, it isn't clear what exactly it is supposed to be. One effort to sidestep this issue comes by saying conceptual truths can be substantial (in the way empirical truths are, presumably). Forget trying to understand what the world's structure is, and instead ask whether the propositions are the right kind to make substantial claims about the world. But what would it mean for necessary truths to be *substantial*?

2. Appeals to the facts and genuine constraints

We can follow Williamson in thinking of substantiality in terms of setting genuine constraints on reality. To set a constraint on reality is to have the world answer to that constraint. To meet the constraint, the world must manifest certain features and not others, as set out by the constraint. And those features it does not manifest must be genuine features that are ruled out by the constraint. If they weren't genuine, I take it that the constraint wouldn't be either, for something can hardly be said to be constrained if there are, in fact, no restrictions placed on it. If there are no real cases outside of the constraint, then there is no state of affair that cannot obtain because of the restriction. And so, it would be strange to talk about a 'constraint', given the lack of

[18] Sider (2011, 194).

meaningful alternatives. I am not constrained by the fact that I cannot be a qwergovitz, if there is no such thing as a qwergovitz (for in that case, there is nothing that I am not or cannot be).

For a proposition to place a genuine constraint on reality, then, is for there to be genuine cases that would refute the relevant proposition if those cases were instantiated. However, in the case of a necessary truth, such alternatives are not just false, but necessarily so – impossible. If necessary truths were substantial, we might say refutation is impossible because the genuine alternatives cannot be instantiated.

We are seldom told what it takes for something to be a 'genuine' case. Indeed, Williamson himself is not especially illuminating here. One might naturally think that a criterion for something's being genuine would be that it is possible. But clearly that cannot be the view here. A minimum requirement, I suggest, is that it is a state of affair that, while impossible, its description does at least carry some content. It is a meaningful proposition even if, in the end, what it tries to depict is impossible. Again, a comparison with empirical propositions might be proffered as helpful. One might say: it is impossible for a human being to jump from the surface of the earth to the surface of the moon. But even though this is impossible, we have still ruled out a genuine case. We can understand what one means by 'a human being jumping from the earth to the moon' even though it is impossible for one to do so. Of course, the kind of impossibility here is physical, whereas the philosopher is interested in one of a higher grade. One might say: philosophers are concerned not just with the world we happen to inhabit, but with the possible worlds we might have inhabited. And so, the philosopher might think they have extended their analogy far enough to make good on what it takes to rule out genuine cases. The analogy would only hold up, however, if we were satisfied that we can understand what exactly is being ruled out by a given necessary proposition.

The trouble is that it is hard to see how such a requirement would be met. Insofar as necessary propositions go some way to setting out the limits of the concepts contained within them, any breach of such a limit is going to amount to a change of subject (at best). To 'deny' that a square has four sides is to misunderstand the meaning of 'square'. What would the genuine case look like here? 'A five-sided square' – but what would that be? A pentagon? If you were to insist that 'no, on the contrary, I am not speaking of the pentagon but rather the five-sided square', I would likely concede that I am no longer following you. Insofar as you insist that you are speaking as I do, you cannot meaningfully talk of five-sided squares. The further you insist, the more likely I am to be lost. Are you using the wrong word? Are you suggesting that we alter the concept 'square' to incorporate pentagons too? Strange though that might be (and without

obvious advantages), at least I could cast your denial in terms of rejecting a certain way of representing things in favour of another. In other words, I could make intelligible a remark that before seemed like plain nonsense. But the price for making it intelligible is for it to no longer amount to a genuine alternative to squares having four sides. It would be to invent a new concept, with its own corresponding set of necessary propositions. One of which is, perhaps, 'a square must have four or five sides', where 'square' is being used in a different sense.

Now, Williamson does consider this argument, albeit briefly. Briefly, because he thinks it obvious that the argument is (viciously) circular. Indeed, he thinks there is a problem with the argument even in the case of logical truths such as 'All furze is furze'.

> To complain that 'Not all furze is furze' does not express a genuine case is to argue in a circle. For it is to assume that a genuine constraint must exclude some logically consistent case. Since substantiality was being understood to consist in imposing a genuine constraint, that is tantamount to assuming that no logical truth is substantial, the very point at issue. (Williamson 2007, 65)

Williamson's charge is effectively that the argument I have given above begs the important question. For it is as if to say: a genuine case is one that adheres to the limits of our concepts and, therefore, no genuine case can be ruled out by those limits because said cases exist only inside of them. But Williamson's characterisation of the argument is mistaken. As Glock remarks in his review of *The Philosophy of Philosophy*, there is an independent ground for ruling out Williamson's allegedly 'genuine' cases, namely their unintelligibility.[19] Thus, we might reply to Williamson that the problem with these cases counting as 'genuine' is that they are incomprehensible, and that their standing outside the limits of our recognised concepts is a plausible reason why. As I argued above, the difficulty is in understanding what one means by one's words when they diverge from our usual meanings. And, indeed, this would be the explanation I would give if I were asked why I found it unintelligible, namely, that so far I had been given no indication of how to understand what had been said. In this vein, it is noteworthy that Williamson does not himself explain what case he has in mind for 'not all furze is furze'.[20] Perhaps it is too obvious for Williamson to feel any need to expand on what would count as furze not being furze. But I must confess that, like Glock, I find this simply incomprehensible,

[19] Glock (2010, 344).

[20] Williamson is not alone in making such claims without further explanation. Sider (2011, 101) likewise does regarding 'it is either raining or not raining'. He thinks it is just obvious that it rules out a genuine case, namely a world in which it is not either raining or not raining. But to what extent one is impressed by that defence will depend, I suspect, on one's own philosophical prejudices.

taking the words to mean what they ordinarily do. How are we to understand someone claiming to both apply and not apply the concept simultaneously? Relatedly, Glock rightly argues elsewhere that one does not successfully conceive of certain logical laws as failing to hold simply by refusing to assign truth values to statements like 'p = p'.[21] It is one thing to claim a form of words might not be in use, another to think there is a coherent description of cases in which something is not identical with itself.

Williamson's complaint appears to be that we judge cases to be ingenuine merely on the basis that they stand outside our laws of logic or linguistic convention, thereby assuming those principles cannot represent genuine constraints. But that obscures the mechanics of the argument. It presents it as though we start off by assuming certain cases are ingenuine. No such assumption need be made. Instead, we can work from the meanings of the words involved and determine whether a given sentence says anything at all. One needn't start by assuming certain cases are illegitimate but can instead investigate whether an expression carries meaning in our language (whether those words really do depict a 'case' in the first place). This is done by consulting our linguistic practice and the norms that constitute it.[22]

Of course, if we accept that necessary truths set limits on the concepts involved, then we have a general argument for considering statements that purport to deny them as failing to depict genuine cases. For the use of such concepts requires that we accept those very propositions. Given there are no coherent alternatives, a necessary proposition cannot conflict with any other and thus cannot depict reality as being a certain way. Before necessary propositions make contact with reality, alternatives to them are already ruled out. Our language leaves no space for them. What is left are arguments regarding the relevant concepts' use and usefulness.

This (briefly sketched) linguistic understanding of necessity has the advantage of demystifying the link between 'inconceivable' and 'unthinkable', and the impossible. It has for a long time been accepted – in some form or another – that something's being unthinkable is good, sometimes indefeasible, evidence for something's impossibility, where something's being impossible means that it *cannot* exist. Some philosophers have been more cautious than others, suggesting that the link between what we think and what reality is like is far from clear. Thus, Thomas Nagel, in *The View from Nowhere*, leaves open the possibility that in some cases 'our strong convictions of positive inconceivability cannot be relied on as evidence for impossibility'

[21] Glock (2003a, 85).
[22] In his PhD dissertation, Nyseth (2018, 227-233) tries to illustrate how this might work in logic, using conjunction elimination and the law of non-contradiction as case-studies.

but assumes such cases are rare.[23] In the account given above, this link is clarified. Saying something is unthinkable is a quite literal way of explaining the problem. Words are employed without their usual limits that help define what we mean by them, and thus we find ourselves in situations where we are unable to understand the combination of words concerned. Thus, it is not some peculiar property of five-sided squares that they cannot exist (are impossible), rather we have not yet determined what is to count as a 'five-sided square' and so have failed to provide something to think about. To posit their existence would mean no more than whistling in the wind.

3. Appeals to the ineffable

But perhaps the problem here is in assuming that the impossible alternatives are expressible in our language. Perhaps they evade our attempts at representation, thereby eluding our comprehension. Javier Kalhat presents a neat counterargument to the claims I have made. He says of an argument that he, in his own words, charitably attributes to Wittgenstein that 'it trades on a questionable assumption, namely that the impossibility of meaningfully stating a fact entails the latter's impossibility, and hence necessary non-existence.'[24] Nagel describes this idea, which he too attributes to Wittgenstein, as 'an attempt to cut the universe down to size'.[25] The general thrust of this argument is one that has reasonably wide appeal. How can we possibly tell what exists from what we can or cannot say? Kalhat thinks the…

> … explanation appears to assume that what facts there can be is determined by what facts can be meaningfully expressed in our language. This assumption (a Tractarian hangover?) strikes me as implausible. Why should the limits of language determine what nonlinguistic facts there can be? It is perfectly conceivable that there should be facts about reality that are as yet unknown to us, and which our language simply cannot state (and perhaps never will) because it is not expressive enough, e.g., because it does not include the required conceptual connections. Some of the more radical scientific discoveries in the past have in fact called for the development of new vocabularies and new conceptual connections before they could be stated. Could Einstein's theory of Special Relativity have been formulated in ancient Greek? Or string theory in the language of the Babylonians? (Kalhat 2008b, 229)

[23] Nagel (1986, 92).
[24] Kalhat (2008b, 229).
[25] Nagel (1986, 109).

Now, it is certainly true that 'some of the more radical scientific discoveries in the past' have called for conceptual innovation. That is not something I wish to dispute, and Einstein's theory of special relativity does indeed strike me as one of the more plausible cases.[26] But what does it prove? The position under attack appears to hold that if a theory cannot be expressed using our existing vocabulary, then we cannot maintain the genuine possibility that it might be correct, or at least be a better approximation to the truth than our existing one. Kalhat overstates his point here. It is true that some advances in science require conceptual innovation, but that does not necessitate a new vocabulary, i.e. new words to express those concepts. Special relativity can be (and often is, in popular science books, for example) expressed using non-specialist vocabulary that existed prior to Einstein's theory. One can describe concepts without creating special words for them. In any case, I accept what I take to be Kalhat's central point, namely that scientific advance sometimes requires conceptual innovation. A position that denied this would, I think, be absurd. But it is not my position to deny it, and I am highly doubtful that it is Wittgenstein's either.[27]

The crucial point, which Kalhat acknowledges in the paragraph above the one I quoted, is that the 'explanation for the special status of necessary propositions is that these propositions are such that we would not count anything as refuting them'.[28] But this differs significantly from the case of scientific discovery. There it makes sense to talk of refutation and counterexamples, even if none have so far been discovered, and even if the competing theory demands conceptual innovation. To put it bluntly, whatever a scientific theory says we should expect, we can conceive of a world where that, in fact, is not the case. Indeed, many theories are incomplete, insofar as they successfully explain some things, but need supplementing (sometimes replacing altogether) to explain related phenomena. Necessary propositions, however, are different, because any attempt to countenance a counterexample changes the nature of the concepts in question, thereby changing the subject. Whatever it is those counterexamples counter, it is not the initial proposition. We saw this in the case of 'every man is subject to fits of rage'. If that is understood as a necessary claim, then any alleged counterexample is not a counterexample to that proposition but to its empirical counterpart which must employ concepts other than the ones in that necessary proposition, for one defining feature of 'man' (namely, being subject to fits of

[26] Conceptual innovation in science is discussed in chapters three and seven.
[27] Notwithstanding the common misapprehension that Wittgenstein is somehow hostile to science. This prejudice is found in Williams (1981) and Maddy (2014). Rarely is the suggestion based on strong evidence, but it is strongly repudiated by Child (2017), Kuusela (2015) and Mulhall (2009a). Both Kuusela and Mulhall consider it a kind of academic myth, while Child concludes that at most one could say that Wittgenstein thinks it unobvious that there is anything good or desirable about scientific knowledge.
[28] Kalhat (2008b, 228).

rage) is no longer definitional (necessary). Kalhat's mistake is easiest to observe when he later clarifies that…

> I am not claiming here, of course, that there could be an as yet inexpressible fact that contradicted a given necessary proposition. If the proposition is genuinely necessary, there can be no such fact. What I am claiming, rather, is that the impossibility of meaningfully stating a fact does not entail the latter's impossibility, and therefore that the ultimate explanation for the special status of necessary propositions cannot be the impossibility of meaningfully stating any facts that purported to contradict them. (Kalhat 2008b, 229)

The trouble for Kalhat is that claiming nothing would count as a refutation is different from claiming that the impossibility of meaningfully stating a fact entails the fact's impossibility. For there are many facts that I lack the means to communicate, because I have never learnt, and will never create, a whole host of concepts necessary to describe them. Concepts grant us access to the facts, give us a means by which to think about and express them. But such facts are not rendered impossible (or false) by my or anyone else's failure to think or talk about them. Kalhat's discussion centres on thus far untouched logical space. But the issue is more local, namely that there is no such thing as stating a fact that refutes a necessary proposition (of course, there is no extralinguistic fact to confirm them either). While the impossibility of meaningfully stating a fact does not entail the impossibility (or non-existence) of such a fact, the unintelligibility of 'refutation' in the case of necessary propositions does mean that we have – as yet – failed to make sense of these supposed 'facts', i.e. those facts capable of refuting necessary propositions. And whilst there is no sense to them, they are not items to be considered. They are nothing whatsoever. The reason such facts are intelligible in the scientific case is that we understand what refutation looks like there – we can make sense of it. The problem, then, is not with facts that cannot be meaningfully stated in our language, but in the notion that there is such a thing as refuting a necessary proposition.

A square without four sides is unintelligible precisely because we understand that in not having four sides it would not be a square. It is not that something couldn't have different features, only that in so having them it would not be what we count as a 'square'. The problem comes when we think that something intelligible has thereby been ruled out, that necessary propositions are so because their potential refutations are, in fact, impossible. My claim is not: there are no facts which cannot be expressed in our existing vocabulary. In denying that refutation is meaningful when it comes to necessary propositions, I am not ruling out certain

facts, for the entire point is that we cannot meaningfully speak of *refutation* here, and so not of the *refuting* facts either. What we learn is not that there is no fact corresponding to the 'description', but that nothing is said by the expression assumed to state the denial of a necessary truth. There is no coherent description to which a fact might correspond.

This is not a matter of opening fresh logical space with new concepts, but of finding a way (within the remits of our existing concepts) to make sense of the idea that there are genuine cases that necessary propositions rule out. But if all such cases remain unintelligible to us (for reasons we can explain), why would we continue to think there are some? Do we have anything more than empty words when considering these 'alternative cases'? Michael Forster raises a related issue against Barry Stroud's interpretation of Wittgenstein's 'diversity thesis'.[29] Wittgenstein famously argued for the possibility of conceptual schemes radically different from our own. Stroud suggests that although Wittgenstein argued for this through examples, these examples are in the end unintelligible.[30] Nevertheless, Stroud argues, they succeed in showing us that such radical departures from our own scheme are possible. The question of course is this: how could one succeed in showing the overall possibility of radical diversion by providing a set of unintelligible examples? If I make an argument on the back of unintelligible cases, you will rightly think that I have failed to make my case. My case cannot rest on examples that are unintelligible, for those examples are thereby void of any content that could be used as support.

In this section, questions have been raised about the limits of meaningful discourse and what lies beyond them. But one might legitimately wonder: if the notion that necessary propositions place genuine constraints on the world was rejected on the basis that those propositions fail to exclude genuine alternatives, what sense is there in talking of them as limit-setting at all? It would be as if these limits, as Stephen Mulhall writes when characterising Bernard Williams's construal of Wittgenstein as a transcendental idealist,[31] 'fenced us in, keeping us out of a domain beyond the domain whose limits they stake out, as if nonsense were a peculiar kind of sense, or as if there were something we cannot do here.'[32]

Mulhall's caution can, however, be acknowledged without abandoning the talk of limits. Within the limits of conceptual sense lies meaningful discourse, on the other side lies utterances that contain elements for which no meaning has been assigned. It is not that *something* or some *thought* lies on the other side. It is not that I cannot conceive of *the object* a five-sided square, but,

[29] Forster (2004, 158-9, 167-8).
[30] Stroud (1965).
[31] Williams (1981).
[32] Mulhall (2009a, 339).

rather, that these words do not represent something to think about. For this reason, such limits are not to be understood as limitations.[33] When I recognise you as having breached the bounds of discourse, I should resist the idea that I cannot understand *what* you are saying, for the very problem is that, outside of meaningful discourse, you cannot say anything at all.[34]

This may still strike the reader as odd. Have we not recast the limits so we are able to mark precisely what can and cannot be thought and said? Is that not what is done when we effectively exclude 'five-sided square' from our language? Not quite. The point is that some combinations of words represent possibilities, while others do not. In identifying those that do not, I come no closer to making sense of what lies beyond such limits or what those words *really* represent. I merely exclude an empty combination of words.[35]

4. Competing conceptual truths

There is another objection to the argument I have set out which I should address. Doesn't its conclusion rule things out that philosophers are well acquainted with? Isn't philosophy characterised by the attempt to discover what something *really* is? So, political philosophers have been concerned with what liberty *really* is and epistemologists likewise with knowledge. The thought is that philosophers do not merely expound variants of the same concept but rather try to trace the contours of the real thing. Can philosophers be wrong about what they are doing when they enter such arguments?

This latter point is one that Williamson develops when discussing the epistemological conception of analyticity. Williamson imagines a dispute between two people, Peter and Stephen, over the sentence 'every vixen is a vixen'. Williamson suggests that Peter could fully understand 'every vixen is a vixen' while denying its truth in cases where there are no vixens, by regarding universal quantification as existentially committing.

To be clear, the question is not whether Peter and Stephen understand the meanings of *their* words (they surely do), but whether their meanings are shared between them (and, perhaps,

[33] See Moore (2012, 135-36, 214-15, 234, 244-45) for further discussion of this important distinction.
[34] Because there is nothing to be said. This is akin to not seeing something because it isn't there, rather than failing to see something because one is blind. Within Wittgenstein scholarship, meaninglessness (or nonsense) has become a topic of fierce debate – especially between the so-called 'New Wittgensteinians' and more 'orthodox' interpreters. See, for example, Conant (2001, 2004), Crary and Read (eds.) (2000), Dain (2008), Glock (2015), Hacker (2003) and Schönbaumsfeld (2010).
[35] See Sullivan (2003, 209-11), Moore (2013, 247), and PI, §500. These issues are treated more fully in the final chapter.

the wider population).[36] For Williamson's argument to pose a serious challenge to my account, he would have to show that such an understanding of universal quantification did not indicate a diversion in meaning with respect to 'every vixen is a vixen'. The matter concerns what is involved in making a claim of this kind and specifically whether it commits us to saying at least one such entity exists. But what is said by a series of words and what we are thereby committed to is precisely what word-meaning is supposed to explain. Insofar as we run on Peter's understanding, we can readily concede that 'every vixen is a vixen' may express something false. But, if we take 'every' with our usual understanding, it is a conceptual truth. This is a clear example of two distinct understandings of the same word. What other explanation is there for the fact that 'every vixen is a vixen' might express something false for Peter, but true for me?[37] Moreover, what explains the fact that I can readily concede that what Peter says is correct, so long as I accept his understanding of 'every', while – at the same time – accept 'every vixen is a vixen' as spoken from my mouth is necessarily true? The only explanation, I submit, is that we are talking in different terms.

In defence of his position, Williamson assures us that 'Peter and Stephen are emphatic that they intend their words to be understood as words of our common language, with their standard English senses. They are not making unilateral declarations of linguistic independence… Each of them believes that his semantic theory is correct for English as spoken by others'.[38] Peter and Stephen are native speakers who learned English in the normal way. It was only in adulthood that they acquired their strange views. He concludes that 'to stop our logical debate with Peter and Stephen in order to explain to them what the word 'every' means in English would be irrelevant and gratuitously patronizing'.[39]

But what support do these imagined chronologies give to the argument? Surely the point here is the one that Amie Thomasson makes with respect to ontology. We can grant that 'the disputants all *think* they are really disagreeing – it's not [their] sincerity… that's in question, but rather their judgement'.[40] Williamson's characters (who happen to have the same understanding of their dispute as Williamson himself) provide no better evidence for his view than would bold pronouncements made by Williamson in his own voice.[41]

---

[36] See Warren (2022, 51).
[37] Chapter eight's discussion touches on this.
[38] Williamson (2007, 89-91).
[39] *Ibid*, 91.
[40] Thomasson (2009, 445).
[41] See Schroeder (2009, 85-91) and Flanagan (2013) for further criticism of Williamson's argument from disagreement.

English speakers might be unlikely to accept that universal quantification is existentially committing. But that is really beside the point. There will be a fact of the matter as to how each of us speaks, and the findings will likely be similar for each of us (though we should not ignore the possibility of our own idiosyncrasies). Those facts risk distracting us from the central point, however, namely that these disagreements are the result of different senses of the same word. Genuine disagreements may well break out between us about which is correct, or more common, in English, or which we ought to employ. But while it might be false to say most speakers employ Peter's concept 'every', that does nothing to detract from the relevant conceptual truths that hold of his practice. It would just mean the concept they determine is not presently in use. Thus, resolving our disagreement with Peter need not be gratuitously patronising, and exactly how it goes will depend on the kind of claim he is making. Is it about how we speak, or how we should?

How might an argument of the second kind go? One suggestion would be that we should speak using those concepts that mirror the world's structure. The aim of this chapter has been to challenge the intelligibility of such a notion, but perhaps now is a good time to revisit it with an example I mentioned at the start of this section.

It has been argued that liberty is a kind of formal freedom, whereby one is free from external constraints. The idea is that to be free is to have the opportunity to pursue one's goals, without the threat of oppressive laws or violence. Critics of this conception suggest that constraints so understood are too narrow. They claim that central to the concept 'liberty' is the *ability* to pursue one's goals and that existing without the threat of oppressive laws and violence is not sufficient for such an ability. A lack of resources, for instance, can be just as significant a constraint as an oppressive law.

What would it mean here to say that one was the correct concept, or that one tracked closer to the truth than the other? The claim cannot merely regard how we do, in fact, think about freedom, for we have already seen that such disputes are perfectly possible on my account. The inclination is to say: we are not interested in what we think about liberty, but in what liberty *is*! The idea is that our competing conceptions are each approximations of the real thing. To say *that* we need to think of a shared plain on which these understandings of liberty are to compete. But the only time they compete is when it is a question of which of the two is instantiated in a *particular* case (in some cases, both might be or neither). This covers both questions about whether the concept applies to a given context and whether it is employed in a particular language. Yet such cases cannot provide what we need. The question is not whether one

conception is instantiated somewhere but rather which of them mirrors the categorial structure of reality.

However, the notion that there exist absolutely correct concepts requires explanation. One risk is that it ends up asking what *really* counts as liberty while expecting the answer to avoid any discussion of what we mean by 'liberty' (for it is precisely those conditions that are up for grabs). It is as if our talk of liberty, though on the face of it constituted by application conditions we are inclined to assent to, does instead have some independent, true meaning, waiting to be excavated. The reason this is problematic is simple: once we remove our (pre)conceptions of what liberty amounts to, we begin to lose grip of the very thing that was supposedly under investigation. We eventually have no conditions for the application of 'liberty', having removed all content from the term. The departure from those conditions, and hence the word's meaning renders the concept void. We no longer have something to discover. A charmed ideal becomes an empty promise.[42] How could I aim to discover truths about something, indeed how could I even think about it, before knowing what a phenomenon of that kind was? Any analogy with empirical investigations breaks down here, for even when we don't know the nature of a given substance, say, we are still agreed on which substance we want to know about (it might be the one we are pointing at, for example). And it is that substance which must be investigated for us to gather knowledge about its nature. One might (albeit imperfectly) capture the difference by suggesting that philosophers seek the object's identity (its essence), while the empirical investigation concerns natural facts about its constitution.[43]

The metaphysicist imagines an essence or form lurking behind our usual understandings of 'liberty' that might come to justify one of our existing conceptions or one we might later invent (discover?). This requires 'competing' concepts to be aimed at some shared target, for the 'competing' conceptual truths to be about the same thing. Conceptual truths determine concepts and different conceptual truths make for different concepts. On this picture, conceptual truths are also meant to (accurately or not) depict the categorial structure of reality. However, across those statements, there is no reference to that structure, the concepts involved differ, and there

[42] This argument runs against Papineau (2009, 19), Kornblith (2002, 1-2), and Millikan (1983, 73), who all talk in terms of learning about the nature of, or real, knowledge, meaning (etc.) rather than the concepts thereof. It hints at a reply to the familiar charge that linguistic accounts *merely* tell us about concepts, not the objects we are interested in. These things are not as easily separated as critics suggest, for we use our concepts to think about those objects. A concept represents the object in thought, and a concept's meaning determines the objects that are picked out by, and fall under, it. Thus, Papineau's suggestion, that Gettier (1963) shows not just that our concept 'knowledge' imposes a requirement of non-accidentality, but that the real thing satisfies this requirement, is trivial. Our concept 'knowledge' determines the type of thing we are willing to call 'knowledge' and so it is no wonder that the requirement is satisfied. This objection will be dealt with more fully in chapter four.

[43] Many think Kripke's work has brought these together. I discuss this in chapter three.

are no coherent descriptions that deny the propositions therein expressed (because using the relevant concepts involves our accepting those propositions). It is, therefore, neither clear that they make competing claims, nor that they answer to reality. The contrary view thus effectively postulates a super-empirical structure to perform a role in respect of propositions that appear to leave little room for it. The problem is less that they postulate something that is inherently strange (what exactly is an objective categorial structure of reality?), more that the purpose of this speculation isn't there to be served.

Various conceptions of liberty exist, and we often argue over which we should employ. Indeed, it is unnecessary to exclude the kinds of proclamations that got us here, such as 'but that isn't *really* freedom!'. For the point might be: you dress your conception up as a satisfactory account of something of grave importance, yet it fails to account for the reasons *why* it is important, because it is, perhaps, so narrow that worrisome restrictions remain consistent with freedom. We might point towards the significant similarities between the cases where one is classed as free, the other unfree, and argue that whatever the differences, the similarities are more salient.[44] Or I might introduce a new way of seeing things to tease out distinctions that are unacknowledged by existing concepts, just as a biologist might when categorising birds by species. Another possibility is that 'liberty' and 'freedom' signify something important in our lives and the struggle is to codify their uses in different circumstances. So, we might bring cases forward and say things like, 'but here we wouldn't say they are free!'. The point is not to decide between different possibilities now, but to study the discourse in question when it arises. I neither give reasons for saying only one of these possibilities is correct, nor that they are the only possibilities. Instead, I am giving a flavour of how we might come to understand our existing discourse. I make no suggestion that these possibilities cannot be present in a single dispute. Rather, I think there is every reason to expect them to overlap, often without the speakers acknowledging this fact. On occasion, reflecting on the nature of one's disagreement may help resolve the issue.[45] But will it always be possible to separate the different possibilities? I am not so sure. For arguments of the first kind need not only be used in support of using one concept over another. Instead, they might expose that someone's description of our concept 'liberty' fails to account for its purpose, thereby failing in its attempt to characterise an existing (shared) concept. If we were to agree that one purpose of our concept 'liberty' is to capture the difference

[44] Given the example concerns 'freedom', these arguments are deeply political. But similar arguments are also possible with concepts that are less politically charged. They are at the heart of the conception of metaphysics developed by Moore (2012).
[45] See Chalmers (2011a) for a discussion of similar issues.

between prisoners and free citizens, then a characterisation that excludes those undergoing punishment from being considered unfree clearly fails to serve its purpose.[46]

5. Conclusion

In *The Philosophy of Philosophy*, Williamson discusses 'barristers are lawyers' and concedes that the meaning of that sentence is sufficient for truth, by which he means that 'necessarily, in any context any sentence with that meaning is true.'[47] Far from this bringing him closer to my conclusion, this condition barely states more than that the sentence expresses a necessary truth, as Williamson notes.

> Consider any non-indexical sentence *s* that expresses a necessarily true proposition. Necessarily, in any context, any sentence with the actual meaning of *s* expresses that necessary truth and is therefore true. […] In that sense, it is true in virtue of meaning. But how little has been achieved in so classifying it! Nothing has been done to rule out the hypothesis that it expresses a profound metaphysical necessity about the nature of the world, knowable if at all only through arduous *a posteriori* investigation, for instance. No reason has been provided to regard *s* as "merely verbal" or "insubstantial" in a pretheoretic sense, unless one already had independent reason to regard all necessities as merely verbal or insubstantial. (Williamson 2007, 61)

If the arguments of this chapter are correct, we do have independent reason to regard necessities as insubstantial. Those arguments have also begun to sketch a more robust sense in which necessary truth might be accounted for by language. But whatever reasons I have provided, one will likely have many more questions. Indeed, one of Williamson's other examples does itself raise a serious one. When expanding on where else meaning might be sufficient for truth, he introduces 'water contains H_2O'. He considers this proposition's meaning to be sufficient for truth also.[48] But while the prospect of 'barristers are lawyers' expressing a substantial truth about the world will strike many as absurd, the chemical composition of water seems an obvious example of a substantial truth. And, ever since Kripke,[49] it has been widely accepted that such truths are *a posteriori* necessities. Here's David Braddon-Mitchell:

---

[46] Cf. PI §564. Part III explores this in more detail and makes explicit its connections with genealogy. One might find Haslanger (2012, 222-225) helpful in thinking about how to answer, 'what is x?' questions.
[47] Williamson (2007, 60).
[48] *Ibid.*
[49] Kripke (1980).

> Famously, we all began to think that everything we believed about something— say, water— could turn out to be false. And despite this, we could still have thoughts about water, and our words could still refer to water. The trick was to see that not only were all the essential properties of water ones that could be discovered *a posteriori*, but the nature of meaning and reference itself was given by the best scientific theory. The post-Kripkean orthodoxy became that matters of reference were doubly *a posteriori*. (Braddon-Mitchell 2008, 24)

This might spell trouble for the arguments of this chapter, for how could it be that truths we know through experience fail to be substantial? This supposed problem will be addressed below, as will many others that would no doubt be raised in opposition to the arguments made thus far.

(2) The Normative Account of Necessity

What follows is a sketch of the account of necessary truth I will defend. A rough picture is available from chapter one, but a more thorough presentation is necessary before answering specific objections. Further details will be rendered explicit in my answers to those objections. At no point will I give a history of normative accounts of necessity, but clearly their roots are to be found in the first half of the twentieth century. While my account follows (an interpretation of) the later philosophy of Wittgenstein, similar views were held by members of the Vienna Circle and can be found in the works of Carnap, for example.[50]

The normative account of necessity considers necessary truths to be insubstantial. They are different in kind from empirical truths, where that difference goes beyond the distinction between necessity and contingency. Necessary truths do not describe the fundamental structure of the world, but are rather expressive of conceptual norms, norms of representation, which guide and govern the way we make sense of and represent the world. Such norms do not have to be acknowledged and appealed to throughout our linguistic practice, but they do nonetheless express the regularities with which we speak, and they provide standards of correctness that we can appeal to should there be a misunderstanding.

Conceptual norms are to be understood as partly constitutive of the terms contained within the statement of them. This idea should be familiar from chapter one. 'All bachelors are unmarried men' determines what we count as 'bachelors', namely unmarried men. What accounts for the necessity in such a proposition is the relation between the norm and the concepts it (partly) constitutes. The reason a bachelor cannot, in any possible world (which is to say, in any describable world), be married is that what we count as 'bachelors' are only those men who are unmarried. There is no conceptual space for married bachelors because the norms that constitute our concept 'bachelor' dictate they must be unmarried. To put it bluntly, there is no such thing as a 'married bachelor'. If one 'denies' the statement of a conceptual norm, this will betoken a misunderstanding, or a rejection of a certain mode of representation, rather than a disputation of the relevant facts. If there is no misunderstanding, then the 'denial' must employ

[50] Non-descriptive approaches to modality have a rich history. They can be found – to varying degrees – in Ayer (1936), Carnap (1947), Schlick (1918), Waismann (1965) and Wittgenstein (1953 [2009]). More recently, their popularity has waned, though Thomasson (2020a) proposes one such account, as do many Wittgensteinian philosophers, such as Baker and Hacker (2014). Other examples of sympathetic treatment more recently can be found in Nyseth (2021), Sidelle (2009), Topey (2019), and Warren (2015a; 2017; 2020; 2022). Greater detail about the history of such approaches can be found in Coffa (1991), Hacker (1996a), Nyseth (2018), and the first chapter of Thomasson (2020a). Tripodi (2020) charts the decline of the Wittgensteinian tradition. A history of the Vienna Circle is provided by Edmonds (2020).

some concepts that are different from those contained within the necessary proposition. For to mean 'bachelor' as we do is to use it in accordance with its meaning (i.e., in accordance with the rule that constitutes its meaning, a rule that is expressed by the relevant 'necessary truth'). 'A bachelor is an unmarried man' can, therefore, only be 'denied' if one is using some of those words without their standard meanings. That is, if one is using concepts in that statement that are not ordinarily expressed using those words. But if one employs different concepts, then one does not successfully state the denial of the necessary proposition. One simply uses the same words to talk of something else. It might be helpful here to invoke a distinction made by Adrian Moore, between *denying* a proposition and *rejecting* one. The thought is that rules of representation may be rejected, but never denied.

> To reject a proposition is to decline to think in such terms: it is to repudiate some or all of the very concepts involved in the proposition. To deny a proposition, by contrast, *is* to think in such terms, but to count the proposition false.[51] (Moore 2012, 593)

So far, my explanation has focused on linguistic necessities. That is, on necessities whose basic expression is in natural language. And they will be my focus going forward. But some necessary truths are not straightforwardly linguistic. What to say of the truths of logic and mathematics, for example? While both bear on the meaning of signs and words, it would be overly simplistic to treat them in precisely the same fashion. Logical rules constitute methods of reasoning. 'Logical inference' is determined by what we take to be tautologies, such as modus ponens, modus tollens and contraposition. We might follow Peter Hacker's model for understanding this, who uses negation as his example:

> … the rule that 'p' follows from '~ ~ p' is a part of the grammar of negation. This rule is not happily expressed by the tautology '~ ~ p ≡ p', which says nothing at all. But one might say that the sentence '~ ~ p ≡ p' *is a tautology* does express this rule of inference. (Baker and Hacker 2014, 314)

Elementary mathematical rules can similarly be said to determine what we mean by mathematical terms, such as 'five', 'multiplication' and 'equinumerous', but their primary role is one of constituting a system for calculating quantities, and hence a system that makes possible certain transformations of empirical propositions about quantities.[52] Those rules determine what it is, for example, to successfully multiply quantities together. Of course, knowing the rules of

[51] As Moore notes, essentially the same distinction can be found in Harman (1967, 134).
[52] Cf. Schroeder (2014).

multiplication does not guarantee that we will always arrive at the correct answer. But we will only be following those rules should we get the correct answer.

Necessary truths, then, express, reflect or state norms that guide our language, logic, and mathematics. Those norms constitute the relevant concepts of those domains. Necessary truths may not be obvious, and may not be used in teaching, but it should be possible for them to be derived from linguistic, logical, or mathematical norms. For instance, we would not ordinarily say that '685 × 5 = 3425' is a rule. But we could say that its unassailable truth rests on norms that govern the concepts contained in that statement which are found to involve our treating it so.[53] To count '685 × 5 = 3425' as false requires some alteration of our rules for the use of the signs involved. The statement represents some of the limits of the concepts used in that statement when combined. In this way, we might count it as expressing a conceptual proposition even though it would be strange to think of it as a rule or norm. One must arrive at that answer to be following the relevant rules and using the concepts correctly. There is nothing more to the conceptual proposition's truth than the fact that the rules which constitute the concepts involved require us to accept it. Our accepting it is partly constitutive of our following those rules (or using those concepts).

Putting it this way enables us to draw a distinction between rules and their logical consequence.[54] Both represent patterns of use we accept as correct, but what we call rules or norms (be they linguistic, logical, or mathematical) are typically expressions we use in corrections, explanations, and justifications; while their consequences are statements that logically follow from those rules. Severin Schroeder's gloss on the distinction between rules and their logical consequences is helpful:

> … it may be called a rule of chess that the white queen's bishop is placed on a black square, but one hesitates to call it a rule that the white queen's bishop can never move to

[53] This is an adapted version of something Dummett (1959, 328-9) says. It is adapted because Dummett is there discussing conventionalism, and the view that these truths rest on our having adopted the relevant norms. But whether we have adopted them or not, the relevant concepts are still constituted by these norms. He suggests the former view belonged to the logical positivists, but that Wittgenstein held an even more radically conventionalist view: that every necessary truth was an expression of an explicit convention, by which he meant a decision has to be made to adopt every single necessary proposition as a convention. However, Diamond (1991, 243-66) convincingly argues against Dummett's interpretation, and pays careful attention to Wittgenstein's thought that linguistic conventions are partly constitutive of the meaning of the terms that are used in the statement of those conventions. *Pace* Dummett's Wittgenstein, it is not true that whatever we choose to say is right, is right, for if we were to simply decide tomorrow that some conclusion followed from a set of premises when it patently did not, we would not be inferring (that is, we would not be doing what we currently call 'inferring'). There are standards that must be met for something to count as a logical inference. Conventions, so understood, are not like informal rules for dinner party guests but rather the rules for what counts as a dinner party. The latter is not violated by eating with one's mouth open but by the host refusing to serve dinner.

[54] See Schroeder (2009; 2017) and Thomasson (2020a) for a similar distinction.

> b8. And one would certainly not call it a rule of chess that in the position White: Kc3, Qa8, Be4; Black: Ka1, Ba2, White can mate in three moves. (Schroeder 2009, 104)

To say something is a necessary truth is to say that it expresses or reflects a conceptual norm. 'Vixens are necessarily female' reflects the norm which requires that 'vixen' be applied only to female foxes. Something is not a vixen if it is male (for vixens are partly constituted by their being female), and so any concept that can be applied to non-females would be – in this way – different from our concept 'vixen'. Furthermore, that we are taken to saying 'vixens are female' is necessarily true, or true in all circumstances, can be explained. For if we use 'vixens are female' to express what we usually do with those words (assuming, therefore, that the same concepts are being employed), then the statement must be true, for those concepts are constituted by a norm that determines as much. Norms define relations between concepts, and statements of those norms are counted as true because they correctly reflect those relations.[55]

Notice that this is quite distinct from related, contingent statements, such as '"vixen" means "female fox"' which would generally be understood as an empirical claim about what 'vixen' means in English. And notice too that these claims can be expressed by object-level statements like 'vixens are foxes'. Indeed, often enough when we say of a rule that it is true or false, we will be referring to the empirical linguistic proposition. It is by no means necessary that 'vixen' means 'female fox', however, for 'vixen' could have just as easily been given a different meaning. This does not make statements of meaning irrelevant to questions of necessity, as they can be used to express norms of representation too. Once I know that 'vixen' means 'female fox', this informs me of the concepts those words express and the norms that constitute the concepts expressed. Concepts are, after all, delineated by word-meaning.

There are, therefore, distinct modes of knowledge and error related to the above distinction. One can be mistaken about what a given word means in a certain language, but otherwise understand the rules pertaining to the concept one means to express (as anyone who has attempted to learn foreign languages will know). Conversely, one can be mistaken about what the norms that constitute a concept require. I might, for example, say that 'all unmarried males are bachelors', only to be reminded that it would be quite unusual to count children as bachelors. Or I might correct someone who demands that we decide what to do purely based on the facts, by pointing out that facts are not – in and of themselves – action-guiding. One cannot decide how one should act by consulting the empirical[56] facts alone. One must have an

[55] Cf. Carnap (1963, 916).
[56] So as to accommodate moral realism.

underlying guide (be it a desire, instruction, or some principle that is sensitive to what those facts are). Knowledge of the correct terminology, or whether a given community employs this set of concepts or some other, is empirical. Likewise, we are taught (and therefore learn from experience) the norms that apply. But coming to understand a range of concepts, the norms that constitute them, and what is therefore required for those concepts to be correctly employed can be understood as giving rise to *a priori*, conceptual knowledge. It comes from reflecting on (or being taught) the relations between different concepts. Another way to think about this might be to say that the practical knowledge one gains from learning a language gives rise to the possibility of *a priori* knowledge. One can draw conclusions about (logical as opposed to, say, physical) possibilities from one's conceptual expertise without recourse to experience.[57] Such *a priori* truths are thus accounted for by the fact that they follow from a given set of norms.

The upshot of this account is that our conceptual scheme is – to a degree – arbitrary. We can choose to accept or reject certain norms in favour of others, thereby deciding to employ some concepts over others. This should not be confused with the thought that the norms we reject are thereby rendered false. Nor should it be confused with the idea that we do not have reasons or justifications for some modes of representation. Indeed, given some of the interests we have, it might be irrational to employ one set of conceptual resources over another, better-suited set.[58] While this may be, to a certain extent, a matter of choice, other factors that play a role in determining our conceptual scheme will not be. For example, our cognitive capacities limit which concepts are useful to us. There is, therefore, an equally important sense in which our conceptual repertoire is hardly arbitrary at all. This is not an inconsistency, but merely a reflection of the different senses in which we might say something is arbitrary.[59] It is a crucial aspect of the normative account that necessary truths are understood to be insubstantial, meaning the necessity of such truths is not accounted for by the way the world is or must be. It is the function of a norm, or rule, that accounts for necessity. This should not, however, be confused with the thought that anything goes, or that there are 'equally good' alternatives to the concepts we have – there might not be.

[57] Cf. Glock (2012b, 74-5; 2017, 84-5).
[58] Part III deals with this more thoroughly.
[59] Cf. *Z*, §358.

(3) *A Posteriori* Necessity: A Case of Substantial Truth?

A posteriori necessities are necessary truths that are discovered via experience. They were made famous by the works of Hilary Putnam and Saul Kripke.[60] A paradigmatic case of this alleged phenomenon comes in the form of natural kinds, such as gold and water. The contention is that natural kinds have an inner nature that is discovered by natural science and those discoveries reveal the real meanings of the terms used to refer to them. The meaning of 'gold', for instance, is given by its atomic number. This acts as a constitutive criterion for gold. Something is not gold if it has an atomic number other than 79. Moreover, it never was gold – even in times when speakers lacked the concept 'atomic number'. Thus, our ancestors who may have identified gold only by its macroscopic features were simply ignorant of the real meaning of their term.

The meaning of a natural kind term is, therefore, determined by its extension. The extension is typically fixed by way of ostensive definition. The meanings of such words are generally said to have been fixed by acts of dubbing (baptisms). Later uses continue to be fixed by these acts if they remain causally connected to the earlier events.[61] While there are stereotypes (or symptoms) of natural kinds (colours, textures, tastes), the meaning of the term is bound by whatever the inner constitution of the substance is. Moreover, those substances are said to have their inner natures necessarily. They are defined by them. There are syntactic and semantic features of word-use that, along with the stereotypes, govern what we might call 'linguistic competence', but the meaning of a word is distinct from such considerations.

This story clearly poses somewhat of a problem for deflationary, linguistic accounts of necessity. Those accounts typically claim that necessary truths are expressions of conceptual norms. They are insubstantial insofar as they do not describe features of reality. Those norms are said to constitute the meanings of the terms contained within them. But if these necessary truths are discovered in perceptual experience, then we do, it seems, have reason to doubt that they are insubstantial. And reason to believe that the world determines, for example, that gold is in essence that substance with atomic number 79 (and that this is the *real* meaning of 'gold').

---

[60] See Putnam (1973; 1975); Kripke (1980). Their views regarding natural kinds are often conflated. But Putnam (1990, 55) has explicitly distanced himself from Kripke, and it is far from clear that Putnam is a natural kind essentialist. Hacking (2007) contains a comparison of Kripke's and Putnam's views on these matters. For the sake of this chapter, however, it should not matter much. Putnam will only feature in the explanation (and later criticism) of the semantic story that he and Kripke, for the most part, share. The central contention of this story is that the inner nature of a term's extension, and the causal relations between speakers' uses, play a crucial role in determining what a word means.

[61] While Kripke and Putnam popularised the notion of *a posteriori* necessities and the causal theory of reference, they can be found earlier in Barcan Marcus (1961). See Smith (1995).

In the first section, I show that an argument of this kind rests on a fundamental misunderstanding of the account Kripke and others develop. It is not necessity that is known by experience. Rather, *a posteriori* necessities rely on *a priori* principles that transform empirical discoveries into necessary truths.

The second section considers William Child's attempt to reconcile a linguistic account with Putnam's semantic externalism. I argue that – in principle – it is possible for a linguistic account to tell a story akin to the one supported by Kripke and Putnam. That is, if Putnam is right that speakers intended to fix their word-meanings by reference to the extensions of those terms, a deflationary tale can be adequately told. Indeed, Putnam – especially in his later work – appears to favour a view that eschews essentialism and metaphysical necessity.[62] However, the 'in principle' is an important qualification, for there are limits to the reconciliation that are left unacknowledged by Child's analysis.

In the final section I describe these limitations. I argue that a total reconciliation would fail to do justice to our past and present lives with words. As well as providing clarity to my overall position, this answers further objections that might be made on the back of *a posteriori* necessities.

1. The necessity of *a posteriori* necessities

Before considering the question of whether natural kinds can be captured by a linguistic account, we should first determine whether *a posteriori* necessity poses a serious threat to such accounts. For if it did, any suggestion that a linguistic account could accommodate natural kinds would be undermined. At most, one could say that such an account has the capacity to capture natural kind *talk*, but such talk would equate to no more than a mere imitation of what really goes on, namely the genuine discovery of *a posteriori* necessities. If one could show that these necessities really are discovered through experience, then we would have a reason to dispute the linguistic account independently of its capacity to capture natural kind *talk*. To recall, the worry is that if we obtain knowledge of certain necessary truths through experience, it seems likely they impose genuine constraints on reality, for it was only through experience of reality that they were discovered.

[62] Putnam (1990, 64-70).

But it doesn't seem that Kripke's account of natural kinds poses this problem: that the relevant truths are *necessarily* true is not learnt via experience. Thus, the notion that necessary truth is discovered through experience is misleading. It is the facts that are discovered, not the fact that they are necessary. This distinction is crucial, and it undermines this attempt to mobilise *a posteriori* necessity against a deflationary account. *A posteriori* necessities have an indispensable *a priori* foundation which accounts for their necessity. This observation leads us back to the central question of whether necessary truths can be considered substantial, as it has not been shown that experience accounts for our knowledge of necessity. The necessity derives from principles that are not themselves discoverable in experience, just as the deflationary alternative would suppose. I do not suggest this is news to proponents of *a posteriori* necessity. On the contrary, many are aware of this and think *a priori* principles are an indispensable part of the story. If they are correct, then despite the initial promise, *a posteriori* necessity doesn't undermine the notion that necessary truths are insubstantial by dint of their being *a posteriori*.

What is discovered by a scientist is the structure of a substance, but nothing the scientist finds can tell them to identify a substance with that structure. There is nothing a scientist can discover which would tell them that *this* or *that* feature is essential. Indeed, while they may have scientific reasons for certain classifications (atomic number was found to be better correlated with chemical properties than was atomic weight), there is nothing in that discovery that suggests this property is the one essential to the substance. It might say that this property is the more causally powerful, or interesting, but an extra step is needed before making the jump to this property being an essential or defining feature of that substance. The step would include some principle along the lines of 'those features found to be "super-explanatory" are essential features that constitute substances' identity'.[63] But that principle is not something to be discovered. No experience or experiment could confirm or falsify it, yet it is essential in accounting for the necessity of '*a posteriori* necessity'. Without such a principle being in place, a discovery of gold's atomic number would not license the move to necessity. From the mere fact that something is a certain way, we cannot derive that it must be.

The claim here does not concern the specific principle that natural kind essentialists might plump for. The point is that the principle plays an indispensable role in any argument for *a posteriori* necessity, and that the relevant principle will never be part of an empirical discovery. Thus, Tuomas Tahko, a supporter of natural kind essentialism, writes:

[63] Cf. Godman, Mallozzi and Papineau (2020).

> The core of the Kripke-Putnam framework of NKE [Natural Kind Essentialism] as it is usually understood is that the combination of an essentialist *a priori* truth about a given natural kind essence and empirical information about the microstructure of that natural kind are needed to establish metaphysically necessary theoretical identity sentences that we are all too familiar with. (Tahko 2015, 801)

This is a standard interpretation of Kripke-inspired natural kind essentialism, shared by its advocates and critics alike.[64] The same interpretation applies to other so-called *a posteriori* necessities. Take 'Hesperus is Phosphorus'. The idea is that we could use both words to refer to the same object without realising. Only later might we learn that Hesperus is, in fact, Phosphorus. In so learning, we are supposed to have discovered another necessary truth through experience. That we can reply in the same way to this example should be unsurprising, given Kripke's point is that the cases are structurally similar. In this case, we realise that the object we refer to as 'Hesperus' is the very same object that we refer to as 'Phosphorus'. But we cannot establish via experience the principle of the necessity of identity. Yet it is this that would guarantee the necessity of 'Hesperus is Phosphorus'.[65] Hence, Peacocke writes:

> … nothing in these examples establishes that necessity can be learned from perceptual experience. Experience is necessary to establish that Tully is Cicero. To move from that identity to the necessity, however, we need the principle of the necessity of identity. That principle is *a priori*, and is not learned from perceptual experience.
>
> […]
>
> The *a posteriori* necessities seem always to result from taking some fundamentally *a priori* necessity, like the necessity of identity, and then inferring from it, together with empirical but non-modal information ("Tully is Cicero"), some modal proposition. We do not seem to find cases of the necessary *a posteriori* that cannot be explained in this fashion. (Peacocke 2005, 742)

Whatever one thinks about *a posteriori* necessity, it poses no special problem for thinking necessary propositions are insubstantial. The 'necessity of identity' may even support their insubstantiality, insofar as it seems best explained by the fact that it makes no sense to say something isn't identical with itself. While there is an idiom to that effect ('he isn't himself

[64] A similar idea can be found in Soames (2006, 293; 2007, 36) and Salmon (2005, 195). Lowe (2007), who criticises the *a posteriori* essentialism of Kripke and Soames, also shares this interpretation. It is worth noting that Tahko (2013) thinks Soames's position flies close to the deflationist one.
[65] See Tahko (2009, 344-5) for more on the *a priori* aspect of these kinds of cases.

today'), there is no literal counterpart. Necessity, then, would be better understood as representing the bounds of sense, rather than a constraint imposed on reality. But even if the natural kind essentialist were correct, and some necessary propositions were substantial, the necessity of their claims still relies on principles that are not discovered in experience and, therefore, such truths do not pose a problem by virtue of their alleged *a posteriori* nature.

2. How a linguistic account might incorporate natural kinds

Nevertheless, one might still think natural kinds pose a problem for linguistic accounts. That is, one might be sympathetic to the kind of semantic story Putnam and Kripke develop and worry that the linguistic account cannot accommodate it. This anxiety can be developed in several ways.

One thought is that an account of natural kinds gives us reason to think that words have a *real* or *true* meaning, waiting to be discovered. Some objects are defined by their inner nature and so it is only once we have discovered that inner nature that we can know the true meaning of those terms. Upon such a discovery, we *must* accept the relevant criteria (provided by an object's inner nature) for the use of that word. Thus, we arrive at substantial necessities, such as 'gold is the substance with atomic number 79'. This form of externalism, it seems, 'implies that the "real" meaning of a word must be discovered by science'.[66] And this might lead us back to the thought that what we discover are real essences. It is worth remarking that natural kinds, as a subset of *a posteriori* necessities, seem to have a peculiar force here, for the notion of discovery is more central to the story. We do discover something about the objects in question. The intuitive force in favour of essentialism is perhaps strengthened further by the common appeal to microstructures – underlying properties that are unobservable to the naked eye.[67] In summary, the worry is that natural kind terms, far from having their meaning constituted by linguistic norms, take their meaning from how things are. Necessity is accounted for not by the role certain propositions have in our language, but by how the world is.

Putting aside more general worries that I will explore in sections below, the linguistic account can provide an adequate response here, if required. A linguistic account of necessity can – in principle – absorb much of what is often said about *a posteriori* necessities. Indeed, Child provides one such response. He distinguishes between two senses of 'criteria':

[66] Glock and Preston (1995, 518).
[67] See Needham (2011) for a critique of microessentialism.

> In the constitutive sense, our criterion for a thing's being gold is our standard of what it takes for something to be gold. In the epistemic sense, our criteria for a thing's being gold are the tests we actually apply in judging whether or not something is gold. (Child 2010, 68)

What this distinction allows for is the possibility that we define gold as that stuff which shares its inner nature with *this* →, while knowing our best test cannot yet determine whether our applications are correct. The hope, presumably, is that our best tests will eventually succeed in becoming reliable guides. The deflationist can hold that a convention was set up to define a sample by its inner nature prior to the discovery of that nature. Thus, Hacker's insistence that 'what a word means is determined by convention, not discovery' can still be satisfied in cases where the discovery of constitutive criteria is made possible.[68] A deflationist can also agree with Kripke, who claims that 'scientific discoveries of species essence do not constitute a "change of meaning"; the possibility of such discoveries was part of the original enterprise.'[69] No change of meaning occurs upon discovery because the initial convention stipulated that it would be defined by that very nature. Hence, while we gain knowledge, the semantics do not change.

For Godman, Mallozzi and Papineau, it is the 'super-explanatory' properties that are necessary, and we identify substances with certain properties *because* they are super-explanatory.[70] Again, one could agree with this without thinking their necessity is a worldly fact. Indeed, our *reason* for establishing norms such that substances are identified with these properties could be because they are super-explanatory. It is perfectly plausible that we draw distinctions according to their worldly significance, but that is not the same as thinking those norms are thereby made true by specific facts about which properties are super-explanatory. Such facts are – in any case – natural and contingent. Nor must we think that a more general principle, such as 'substances are identified by their super-explanatory properties', is some deep, *a priori* truth as opposed to a normative principle that, given our interests, it is expedient in certain circumstances to follow.[71]

Crucial to this picture is the fact that the meaning of the term is still fixed by the relevant norm; for instance, that gold will be defined as whatever inner nature it has.[72] The essence of gold, therefore, is still expressed by norms of description. By following a norm that defines gold by its inner nature, we thereby determine that gold will be defined by the relevant findings of

[68] Hacker (2004, 54).
[69] Kripke (1980, 138).
[70] Godman, Mallozzi, and Papineau (2020).
[71] Chapters seven, nine, and ten explore this in more detail.
[72] Though see below for reasons this might be problematic.

scientific inquiry.[73] The necessity is the result of the norm (or convention) that was set up. Insofar as those norms are not themselves substantial, nor are the necessary truths that are later derived from them. While we discover a substantial truth, namely that the stuff has this inner structure, the discovery does not license the move to necessity. The existing norm hardens the discovery into a rule for the use of a concept. In effect, the norm provides a schema into which empirical discoveries are plugged.[74]

None of this is to say that norms determine that the substance we call 'gold' has atomic number 79, for norms of description do not meddle with that which is described. It is just to say such norms account for our making distinctions along these lines. Our language should not be understood as inventing truths and realities. Nor should it be thought of as excluding certain empirical possibilities. Language allows us to compare our empirical propositions with reality. But it does not determine whether there are, for example, cows in a nearby field. Reality does. Linguistic norms tell us what conditions must be met for the statement to be true. They explain what it would mean for there to be cows in a nearby field. They say: when these criteria are satisfied, the proposition is true. But this no more tells me that the empirical proposition is true than the rule of checkmate tells me that I have been checkmated. (I may not even be playing a game of chess.) To find out if I am checkmated, I must consult the relative positions of the pieces on the chessboard. But there is no such thing as being 'checkmated' without the rules of chess. Linguistic norms do not tell us which empirical statements are true. The world, or reality, determines whether there are cows in a nearby field. An empirical proposition being true or false, therefore, does not belong to our language. The possibility of ascribing truth or falsity to such a proposition, however, does.

Before concluding this section, it is worth recognising that Putnam himself comes to be somewhat sympathetic to a deflationary understanding of *a posteriori* necessities. For example, he describes himself as having presented, in 'Necessity and Possibility',[75] a theory 'which was related to Kripke's, but which was stripped of metaphysical assumptions to the point where Carnap might have accepted it'.[76] It was only later that he realised it is not possible 'to assimilate

[73] Sidelle (2002, 319) takes a similar position to this. See also Sidelle (1989).
[74] So far as other *a posteriori* necessities go, a similar story can be told as was done in the case of natural kinds. Take 'Hesperus is Phosphorus'. The linguistic philosopher can claim that we have set up two expressions to refer to, unbeknownst to us, the very same object. What we learn from experience is that we have named the same object twice. We are then licensed to move from 'there is Hesperus' to 'there is Phosphorus', given our discovery and what we mean by saying something is the same object. If we had meant only to refer to that object in the sky at that time of day (or night), then such a move would not have been licensed. But once these rules are in place, 'Hesperus is Phosphorus' follows trivially from the discovery that each expression refers to the same planet.
[75] Putnam (1983).
[76] Putnam (1990, 64)

[Kripke's] *metaphysical* intuitions to the *linguistic* intuitions that other analytic philosophers talk about'.[77] Indeed, Putnam cites his worries about the notion of 'metaphysical necessity', which developed over time, as the reason for wanting to distance himself from Kripke more than he did in earlier work.[78] And he eventually rejects 'metaphysical necessity' in the context of the kinds of examples being discussed.[79]

3. Against total reconciliation

But while it is – in principle – possible to tell something like the Kripke-Putnam story in a deflationary way, there are significant problems with how it is applied, especially in the case of natural kinds. The application of natural kind semantics is severely limited, and Child, like Putnam and Kripke, doesn't mention this in his discussion. This is not to deny that a perfectly plausible account of these necessities is that they harden an empirical discovery into a linguistic norm. That we only count stuff with atomic number 79 as 'gold' could explain the relevant necessity and, as Glock suggests, why 'appeals to linguistic intuitions or "what we say" in certain counterfactual situations' support claims about *a posteriori* necessities.[80] But that doesn't commit us to many features of the story often told. Neither diachronic meaning invariance nor discovery of word-meaning, for example, are thereby presupposed.

One reason this might be important, aside from warning against the uncritical acceptance of a deflated version of natural kind essentialism, is that the standard story might be naturally resistant to the deflationary move. The presence of a deflationary alternative, while effective at blocking a general objection from the mere existence of *a posteriori* necessities, doesn't thereby make that alternative plausible. For instance, it might be suggested that there is *no* conceptual content in the baptising of objects (and reference is determined directly, or purely causally). This might then be used to dispute that *a posteriori* necessities are in any way semantic. Separately, the notion that we discover otherwise entirely opaque meanings (and definitions) through science might support a more substantial picture of those necessities. That is, it might be understood as a matter of discovering the essential properties of the substances referred to. A critique of the semantic story associated with natural kinds therefore supports the plausibility of a deflationary account.

---

[77] *Ibid.*
[78] *Ibid*, 55.
[79] *Ibid*, 70.
[80] Glock (2002, 239). See also chapter four.

The next three sub-sections tackle the former, 'referentialist' side, while the final three assess the extent to which we *discover* meanings. Combined, they warn against a simple translation of the semantics often associated with natural kinds into a deflationary mould and support a more semantic understanding of *a posteriori* necessities.

3.1 The qua problem

The first thing to note is that at least some conceptual content is required to fix reference. A mere baptism (the pointing of a finger, for example) is not sufficiently determinate. One way of trying to get a direct reference theory off the ground is to say we refer to the things that share the relevant causal relation to the naming ceremony. But objects and causes can be described in many ways. Some conceptual content is necessary to determine which we mean. For instance, if I point to a brown bear's footprint and intimate that I am referring to whatever left that mark, I could be speaking about a whole host of things, including: 'living thing', 'animal', 'mammal', 'bear', 'brown bear', 'female brown bear', and so on. The same would apply if I saw the bear itself and gave it a name. An initial encounter such as this does not fix a reference without further ado. Some additional context is required. This point is made by Wittgenstein in the *Investigations* with respect to ostensive definition.

To be clear, this is not fatal for causal theories in general, if what one means by this is theories that incorporate some causal element. We might define a term by using a description that appeals to a cause. The *animal* or *species of animal* that left the footprint, for example. To what extent this remains a 'causal theory of reference' is a moot point.[81] The conceptual content that is included is all the deflationist needs to make their story a plausible alternative. It is decisive against the objection that referential theories of meaning show *a posteriori* necessities could not be semantic.

3.2 Non-observables and non-referring terms

The failure to discriminate between different entities without conceptual content is exacerbated in cases where terms refer to non-observable entities. While we might be causally connected to non-observables, it is unclear how we could discriminate between causal connections that are

[81] See Bird (2000, 186) for a discussion of this. 'Causal descriptivism' (see, e.g., Kroon 1987) is often the name given to an account that combines causal-historical chains of reference with descriptive conditions.

connections of reference and those that are not.[82] Non-observables might be said to face a problem one step before their observable counterparts, for at least an initial ostensive dubbing provides *some* context as to what is being referred to. Some descriptive content (the role a concept has in a theory, for instance) is necessary to identify the non-observable item.

Alexander Bird rightly identifies that this problem spills over to terms that fail to refer.[83] Simple referential theories of meaning appear doomed to conflate such terms. A classic example would be phlogiston, a substance once thought to be released during combustion.

> There is no phlogiston, and so there can be no causal connection between phlogiston and our use of the term. In which case 'phlogiston' fails to refer. That in itself is the conclusion we should expect. But two potential problems are raised. First, in this respect 'phlogiston' is in the same boat as 'caloric', 'N-ray' and so on, all of which fail to refer. The simple causal theory makes no distinction between them, yet there is surely some sense of 'meaning', close perhaps to intension, in which the meanings of these terms differ. (Bird 2000, 184-5)

Bird goes on to suggest a further problem for this account.

> [Phlogiston] was hypothesized in an explanation of combustion. What actually is involved in instances of combustion is oxygen. Hence there is a danger that because oxygen is causally responsible for combustion, which in turn is causally connected to our use of 'phlogiston', it may be a consequence of the causal view that 'phlogiston' is in fact a name for oxygen. (Bird 2000, 185)

But it is, of course, clear that those in the past were not merely referring to whatever caused combustion by using 'phlogiston'. The reason this is clear comes from the conceptual content that explains how the term was used.

> Because phlogiston theorists held quite specific beliefs about the nature of phlogiston and about its role in the causal processes underlying combustion and calcination, they failed to refer to oxygen. They did not mean to refer broadly to whatever it is that in fact causes combustion and calcination, but to a specific substance which produces those effects in a particular way. (Sankey 1994, 65)

[82] Cf. Bird (2000, 184).
[83] *Ibid.*

It is a distinct question whether someone can refer using the resulting expression without knowing what that conceptual content is. One contention of Kripke and Putnam's causal-historical accounts of reference is that they can. I suspect that is probably right: words generally belong to languages where what they refer to is settled, so someone need not have a complete grasp of the word to use it to refer. But there are limits to this thought that are not obviously respected by causal-historical accounts, such as when an entire community uses a word to talk about something that the causal-historical chain implies is not being referred to. Of primary importance is not the causal-historical chain but the linguistic practice the word is used in and the norms by which it is applied.[84]

3.3 Putnam's cats

An argument of Putnam's is sometimes used to object to the claim that conceptual content is necessary for disambiguation. He argues that the category to which an item belongs is often not a matter of necessity. One example involves our discovering that all things we've called 'cats' are, in fact, robots. He suggests we would not say there were no cats in this scenario, but rather that we had discovered that cats are robots.[85]

Three responses can be made here. Firstly, even if Putnam is correct, this does not rule out that cats are necessarily animals. For the relevant semantic rule might be conditional: if the empirical facts are as we think they are, then cats are necessarily animals. This would provide the reference-fixing conceptual content where the empirical circumstances are as expected. If Putnam is correct, there might be a distinct, implicit conditional rule stipulating what happens if all cats turn out to be robots (or a more general rule that applies to this case). What the conceptual norms are is irrelevant at this stage. The point is that what Putnam says can be made consistent with a linguistic account of necessity.[86] Alternatively, it might be that the fact cats are as we believe them to be fixes the reference of our term in imaginary counterfactuals such as Putnam's, meaning he is wrong about that case. Reference is not only fixed in the relevant circumstances but rather the relevant circumstances partly determine how the reference is fixed *in general* (i.e., for all circumstances, including imaginary ones).[87] On this account, cats must be animals, but had things turned out differently we might not have accepted this norm.

---

[84] See 3.5 and 3.6.
[85] Putnam (1962b).
[86] See Thomasson (2020a, 92-112).
[87] This is similar to what Kripke (1980, 126) says about such cases.

Secondly, even if Putnam is correct about his peculiar case, it does not follow that the lesson generalises. We would not always be willing to accept that we were referring to whatever happens to be impersonating the expected referent. One plausible example would be where, instead of all cats turning out to be robots, it is one species of cat. In this kind of case, we would be more inclined to say there is no such species. They all turned out to be fake. This is not dissimilar to when a new species of something is said to have been discovered only for it to turn out that the referred to animals are really members of an already known species. There we tend to say the 'new species' doesn't exist.

A different example, again suggesting the lesson doesn't generalise, could be taken from the online phenomenon of catfishing, which involves creating a fake persona and deceiving others on social media. It seems to me that both ways of describing this situation are permissible: the person a victim believes they are speaking to – in a sense – doesn't exist, and in another is someone else. But if it turns out there are multiple people engaged in the catfishing, then it seems more likely we would say the created persona doesn't exist, rather than that the name refers to the wider group.

Thomasson makes an important point about this, which she borrows from Imogen Dickie: whether a given description is appropriate may be context-dependent.[88] Here Dickie provides a real-world example. The mathematics papers attributed to Bourbaki were authored by a group not an individual. Dickie suggests that the reaction to this might differ between, say, a biographer of Bourbaki and a mathematician. The biographer is much more likely to declare that Bourbaki doesn't exist, while the mathematician might be quite happy to continue to refer to the group as 'Bourbaki' – in part because the personal history of the authors is not relevant to what the mathematician is doing.

Thus, even if what Putnam says about his robotic cats is correct, the lesson doesn't generalise. It may be indeterminate whether – or to what – a term refers in unusual circumstances. It may be down to choices we make post-discovery to determine it going forward.[89] This would fit well with the critique I am making, for the point is that conceptual content is required to disambiguate reference. As Stewart Shapiro and Craige Roberts observe,

[88] Thomasson (2020a, 101-2); Dickie (2015, 168-9).
[89] Cf. Jackson (1998, 54). See also Shapiro and Roberts (2019, 192-97) and discussions of open texture in chapter five and further discussions of the importance of common contextual conditions in chapter nine.

the potential for indeterminacy to arise runs counter to the suggestion that natural kind terms can fix reference for all circumstances.[90] This theme is developed below.

Finally, Putnam's example – even if it did generalise – does not avoid the qua problem. Even if 'cat' does refer to whatever those things we call 'cats' turn out to be, this doesn't fix the reference any more than before. Is it a particular time slice of the robot? Is it this specific robot, or robots in general? Or just robots that share this appearance? Or some particular configuration of robots? Or just those robots we happened to have specifically referred to in the past? Ostension alone is not fit to determine this.

3.4 Conceptual refinement

There are further difficulties for the stories typically told of natural kinds, however, even when conceptual content is included. One is that some descriptions are not themselves sufficiently determinate. Take 'inner nature', for example. What makes the deflationary story intelligible in cases like 'gold is that substance with atomic number 79' or 'water contains H_2O' is our understanding of chemical elements. But such an understanding was not available in the distant past, so what sense does it make to suppose that before the development of modern chemistry, for example, talk of the 'inner nature' of gold or water could be straightforwardly mapped onto how we discuss it now? In what sense can the mere talk of 'inner nature' determine with any precision its own extension? What are we to make of the idea that those in the distant past really meant what we mean now? 'Inner nature' is open wide to interpretation and would not single out any one thing without further ado. One might like to appeal to the inner nature responsible for something's macroscopic properties as a way of determining the microscopic features that matter. But firstly, which macroscopic properties? And secondly, what if the same macroscopic properties have different microscopic causes? Is it obvious that we would (should?) characterise those substances as distinct? Even if we do characterise them as distinct, how would we determine which of them was being referred to by our ancestors? And what if we found that the same microscopic properties cause significantly different macroscopic features depending on their surrounding environments? And what if they did so at random, independent of those environments? Furthermore, if microstructures were discovered to be in states of constant flux, what then would our distant ancestors have been said to have meant?

[90] Shapiro and Roberts (2019).

To move from hypothetical questions to those we faced: if we are to remain in keeping with the meanings of our ancestors' words, should we categorise elements by atomic weight or number? What should we do about isotopes? Couldn't our ancestors, on this account, be said to have referred to the most common isotope of a given element, rather than the element itself, for example?[91] These questions cannot be asked, let alone answered, prior to a relatively sophisticated understanding of the relevant sciences. There are decisions to be made before 'inner nature' in the mouths of distant ancestors could be made sufficiently determinate. Yet these decisions were not ones that could be made by those ancestors, so their notion of 'inner nature' cannot be said to have *really*, all along picked out the relevant attributes. We cannot anticipate every decision that will need to be made prior to our investigation into the nature of water. We do not know what we will find. Indeed, have we any reason to think it so much as possible that our ancestors meant chemical structure when they talked of 'inner nature'? What would it have been like for them to have meant something else? In present times, it is perfectly natural for us to move straight to the chemical elements water is made up of when discussing its 'inner nature', but that is with a perspective those in the distant past lacked.

Putnam thinks that when Archimedes asserted that something was gold, he was making some claim about its inner nature.[92] I do not know what evidence he has for this claim. In any case, my point here is that even if Putnam is right about Archimedes' intentions, the term 'gold' in Archimedes' mouth would not have been sufficient to determine that gold is that substance with atomic number 79. Decisions are made along the way, where the relevant alternatives to be chosen from are not available prior to significant discovery. For similar reasons, whether Archimedes upon being presented with our discoveries would have agreed to categorise gold in the way that we do[93] is irrelevant to what was meant prior to those discoveries. He is free to

[91] Donnellan (1983) presents a case in which two scientifically minded communities, who are otherwise very similar, come to use 'gold' differently (one uses it as we seem to, namely to refer to the element with atomic number 79, while the other refers to the most common isotope of that element). He thinks this diversion is plausible even under the machinery of reference proposed by Kripke and Putnam. The upshot is supposed to be that nature does not fully determine the extensions of vernacular natural kind terms. Schroeder (2017a, 257) emphasises something similar, when he points out that we still have to treat words in a certain way for them to take on particular meanings, to decide *which* features of a sample to take as defining features, even if we think them natural. Thomasson (2020b) shows how thinking that some concepts are more natural than others or thinking the world – in some sense, and in conjunction with our values and interests – justifies our use of certain concepts, does not entail a commitment to a metaphysical structure that our concepts might be said to answer to. Indeed, as Thomasson points out, the factors influencing our concepts are straightforwardly empirical and natural facts. Metaphysical structure need not enter the picture. Meanwhile 'intuitions' regarding metaphysical structure can be accounted for by explaining the way natural facts influence our conceptual scheme. An important lesson is that while it might be a contingent matter which concepts we have, and while they may not answer to the world in the fashion empirical propositions do, this does not imply there aren't reasons for using our concepts. See Part III.

[92] Putnam (1975, 235).

[93] Ibid, 238.

make the same decisions we do, and for the same reasons, but it makes little sense to suggest they were anticipated in his initial intentions when the questions could not have been asked.[94]

Furthermore, even when a description is sufficiently determinate at one time, it might not be at another. Our concepts over time are refined. What is at one point taken to be an unproblematic, sufficiently determinate, definition may later demand to be made more precise. Joseph LaPorte explores this in *Natural Kinds and Conceptual Change*. He finds that, at the margins, adjustments are often made to empirical concepts. Certain discoveries, for example, can challenge us to refine our categories. If it is indeterminate whether a new discovery fits a certain description, we might choose to refine the concept to give a definitive verdict. When a term is made more precise with respect to some case(s), it undergoes a subtle change in meaning.[95]

> [T]he giant panda is now considered to be a bear. But systematists disagreed about its status from the time it was discovered by Europeans in 1869 until close to the end of the twentieth century. Some systematists said it was a bear, some said it was a raccoon, and some said it belonged to a family of its own. The panda looks like a bear, but it has characteristics that set it apart […] The panda's head has a different shape, its diet is vegetarian, and its hand has an opposable 'thumb.' The panda does not growl or roar as other bears do but rather makes a sound that resembles the bleating of sheep. The panda also does not hibernate as other bears do. […] The panda is like the paradigm bears in some respects but not others. Is there a fact that the panda was a 'bear,' as speakers of 1869 used that term, or a fact that it was not a 'bear'? I do not think so.[96] (LaPorte 2003, 83)

In other words, the panda was a borderline case and it's plausible that there was no fact of the matter as to whether it counted as a 'bear' because the application conditions for 'bear' at the time didn't settle the matter.[97] One might suggest that ancestral connection is a surer way of determining whether the panda was a bear, but similar problems can arise for evolutionary taxonomists. There are intermediates in genealogical trees just as there are in phenotypic taxonomies.[98]

---

[94] For more detailed criticism of Putnam, see Hanfling (1984).
[95] See Part II's discussions of open texture and conceptual change.
[96] For other examples, see LaPorte (2003, 71; 67-8; 96; 101-2; 106).
[97] Dupre (1981; 1993) may also be relevant here; but one should acknowledge the cautionary notes LaPorte (2003, 29-31) strikes about jumping from Dupre's observations to conclusions about natural kinds.
[98] Cf. *ibid*, 83-85.

One might likewise think conceptual refinement could take place with substances in a story not entirely different from Putnam's. A given term might be defined by reference to paradigmatic samples and that which bears certain (perhaps both macroscopic and microscopic) similarities. With little more detail, this is a relatively vague term. But as investigations into different substances advanced, we may have come to define with greater precision what those underlying similarities must be. This would be a form of conceptual refinement.

3.5 (Mis-)Understanding past (and present) language-use

One reaction to this might be to suggest that the meaning of our term is whatever future scientists settle on. Or, if we are dealing with a natural kind essentialist, the meaning of our term is determined by whatever the essence of the kind really is. In other words, the refinements that take place are just the scientists' latest attempts at describing essences. Both responses bring out what is perhaps most strange about this brand of semantic externalism, namely that it implies the meaning of a term can be opaque to all users. This has strange consequences.

For one, what use would such a term have been to our ancestors? If 'water' just meant whatever has the same inner nature as this → stuff, how could 'water' have been applied? They had not the faintest idea of which stuff shared water's nature and which did not. After all, they did not know what water's inner nature was. Given they were fully aware of their ignorance, the word 'water', insofar as its meaning could genuinely be explained by appeal to some (unknown) inner nature, would not have been the least bit useful and its application, to be consistent with that meaning, would have been severely restricted.[99] My point is not that they did not use a word for what we, for the most part, call 'water', but rather that the very fact that they did is good reason to think their word did not *really* mean 'H_2O'. The problem is not that they did not each have the means to check whether everything they applied 'water' to did, in fact, share the same nature, for even if we were to accept H_2O as a constitutive criterion for water now, it does not follow that we test everything we apply the word to. Once we have the constitutive criterion, i.e., once we employ it as a standard of correct use, then we can, of course, make errors in our application of 'water'. We can even be systematic in those errors, if we, for example, are wrong about what the symptoms of H_2O are. But without the constitutive criterion, we cannot have associated symptoms (about which we might be mistaken). For one, we neither know how often

[99] See Hacker (1996a, 252) for more criticism on the possibility of defining a term by ostensive definition and appeal to a 'sameness relation'.

that substance presents these symptoms, nor how many other substances present the same way. More importantly, we have no way of finding this out – hence our application of the term should be severely restricted if its meaning really is determined by some hidden nature.

We cannot track the symptoms (inductive correlations) of *X* before having settled on what *X* is. That would be like asking for a gender breakdown of the candidates that had passed a first-round assessment without having decided who had passed the first round. Of course, we can encounter a set of features that we later come to recognise as symptoms for *X*, but that does not mean prior to the conception of *X* that whatever we called that group of symptoms was our way of meaning '*X*'. For if that set of features are merely symptoms of *X*, *X* must have constitutive criteria separate from them. It is these constitutive criteria which determine the meaning of '*X*', and it is precisely those criteria that I lack when I have only a set of what I later classify as symptoms of *X*. Thus, even if in the past we used '*X'* as the word to capture the set of symptoms, we did not mean the same thing by it as we do now. While our use may have been, for the most part, consistent with the modern-day use of '*X'*, there are nonetheless different standards of correctness across the two cases. For constitutive criteria, in the end, win the day against symptoms. When in the modern-day we can admit mistake upon discovering the symptoms without the constitutive criteria, in the past we could make no such distinction. What we now class as symptoms did once act as constitutive criteria. It is a moot point whether *X* survives this alteration, or whether we are better off thinking of it as something new. I suspect most will be inclined to say that it survives, though this does not mean nothing has changed. The laws of rugby change frequently, but we are not inclined to say a new sport is being played.

At this point, an opponent might argue that the kind of systematic misuse of the term 'water' made by our ancestors was done in full knowledge of what they really meant. If they had meant all along for 'water' to be defined by its inner nature, then what does it matter if they go on to use it to refer to substances with the same macroscopic features as the original sample? 'After all', they might say, 'isn't that precisely what happens in the case of words like "wicked", which means "evil" but is now as commonly used in place of "excellent"'. The key detail missing in this story, however, is how to explain changes in the word 'wicked'. My grandmother might be shocked when my cousin, upon opening her Christmas present, says 'that's wicked!'. But what we would say to explain this to my grandmother is that 'she doesn't mean it's evil but that it's great', with some explanation that this has become a standard use of that word. What my cousin is saying is that the present is great (whether she believes that or not). The word has evolved to mean something other than 'evil' in certain circumstances. Now, what should be our response if my cousin were to say, 'no, "wicked" means evil, I just use it instead of "great"'? Well, to take

little notice. It is not for my cousin to systematically use a word in accordance with one meaning and then say she means something else. Something has gone quite seriously wrong if the meaning of a word is so far divorced from its use that one can regularly use a word to say 'great' without it carrying that meaning. One hardly needs to subscribe to a 'use theory of meaning' to think this absurd. What a word is used to express or convey is not independent of its meaning.

One might hope to mount a counterargument to this kind of explanation by appeal to sarcasm or hidden intentions. The thought would be that in such cases words continue to mean the same thing, while we intend for something else to be communicated by using them. Thus, words can be used to express something distinct from what the words alone mean. The first thing to notice is that I never denied this. I merely suggested the meaning of a word and what it says, or expresses, are not independent of one another. Words must be used in accordance with their meaning, but that does not mean they are only used to express what is meant by those words. For example, to say 'oh, that's great' in a sarcastic manner relies on the meaning of 'great' being different from what I intend to communicate to those attuned to sarcasm. Otherwise, it would not be an instance of sarcasm. Likewise, if my partner believes her friend wishes to speak to her in confidence, she might suggest that I go check on the dinner to save her friend from the embarrassment of having me banished to another room. But, again, that embarrassment is only saved because those words keep their usual meaning. If the phrase 'you should check on dinner' was a literal translation of 'my friend wishes to speak to me privately, please go away', then clearly the desired effect of my partner's discretion would be lost. In other words, nothing I have said need abolish the distinction between speaker meaning, or a speaker's intentions, and linguistic meaning. But no equivalent explanations are forthcoming in the cases of 'wicked' and 'water'. My conclusion in the case of 'water' (or some word like it) was precisely that our ancestors used it in accordance with a meaning other than just 'H_2O'.

For what it's worth, it isn't clear that our word 'water' functions today in the way this example supposes. 'Water' does not simply mean 'H_2O'. It is not H_2O that has a high mineral content and is therefore called 'hard water'. 'H_2O' is the chemical formula for what we call 'ultra-pure water', but we say that water with a low mineral content is soft water. If we were to rely on ultra-pure water (H_2O) for drinking water, it would begin to take electrolytes from our blood, leading to illness and eventually death.[100] Moreover, we call D_2O (an isotope of H_2O with twice the atomic weight) 'heavy water'. In other words, there is a mixture of good candidates for water and simple identity statements don't seem best placed to capture them. After all, we resist

[100] Cf. Schroeder (2017b, 118-19).

counting ice or steam as 'water'.[101] Must water contain H_2O (or an isotope thereof)? Possibly, but perhaps this restriction is limited. In Putnam's famous example of 'twater' on Twin Earth, it is far from obvious that our current concept is determinate on the issue (since it has never arisen as a serious question). It seems to me this is another decision for us to make,[102] not least because a different chemistry appears to apply on Twin Earth for which our vernacular has not been prepared.[103]

The key point of this section, however, is that claiming terms used in the past *really* meant whatever the inner constitution of the referent is can have bizarre consequences for past language use. It may also have strange consequences for *our* language use depending on what is discovered in the future. And we have already seen the consequences it can have on terms that are later found not to refer. That word-meaning is bound to linguistic practice is mistakenly overlooked.

3.6 'Word-meaning'

The broader issue here is the fundamental disagreement about what role word-meaning plays. This often-told semantic story risks divorcing 'meaning' from things we otherwise think it is conceptually wedded to. Crucially, it conceives of meaning as being independent of speakers' desired communicable ends. This is so far away from what we otherwise understand by 'word-meaning' that it requires significant motivation. Even with this motivation, however, it isn't clear what relevance the new concept has to debates about meaning, for they are principally about a phenomenon the new concept cannot explain, namely how we are able to speak to, and understand, each other.

One motivation for Putnam's externalism is that he (rightly) thinks a word does not have meaning simply by my intending you to understand something by it (or by some other mental state). But it does not follow that meanings are principally determined by the objects we use words to talk about. John Preston is right that this opens the door to what is an absurdity, namely semantic scepticism:

> … one valuable lesson from Wittgenstein's later work was that the fact that some 'realist' or 'referential' conceptions of meaning leave open the possibility of semantic scepticism

[101] Cf. Stroll (1989).
[102] Cf. Schroeder (2006, 248-9).
[103] Cf. Kuhn (1990).

> should be regarded as a decisive objection against them. […] Scepticism about claims that we know the way things are often makes sense. But scepticism about *meaning* is ridiculous: the meaning of words and statements in human languages is determined by human activity, and must be humanly and publicly knowable. (Preston 2004, 329)

Preston still wants to include the notion of reference in his conception of 'meaning'. It shouldn't turn out that we don't – in general – know what our words refer to. For that too would mean we were talking about things without knowing we were. But this doesn't rule out that we can discover facts about those referents. There is plenty to be learned about the stuff we refer to (their chemical compositions, for example). The point is that we know what counts as the substance we need to investigate to find this out. How else could we investigate it? If we later make some of those discovered properties defining features of the substance, this will amount to an instance of conceptual refinement.

The externalism under consideration divorces language from its otherwise accepted function, namely intentional communication. If it is possible that all of us – subject experts included – might be ignorant of the true meanings of our terms, it follows that word-meaning is not responsible for our being able to communicate what we want to communicate. As Preston points out, this consequence renders this account – as an account of word-meaning – entirely absurd. One requirement for our concept 'word-meaning' is that it explains how we can speak to, and understand, one another. Language has developed to fulfil this function and word-meaning is a concept designed to explain how scribbles on a page, or the noises we make, can do that. What matters is what the word signifies, not what it may have been used to refer to at its initial baptism (or what others will later agree to use it to refer to). Word-meaning cannot transcend the linguistic practice it depends on.

To be clear, the notion that we might *discover* the meaning of a word is not entirely absurd. For one, learning a new language might be described as discovering meanings. But there are also circumstances where something like the crude externalist story might be right. Here is one (highly artificial) case. In principle, I could stipulate that I will use the word 'adon' to denote animals that have as many legs as there are coins in Charlotte's pocket at a specific time. Thus, there is a sense in which I would come to understand the meaning of my own term upon discovering that there are three coins in Charlotte's pocket at that time. But the point is that what gives 'adon' that meaning is that we use it in that way (it is not independent of the criteria we adopt). Words are human artefacts, and their meanings depend on how we behave. Far from requiring a discovery of adon's deepest essence, I can be said to have learnt how many coins are

in Charlotte's pocket and acknowledge that, owing to the norm I follow, this entails that adons are three-legged animals. This story is, of course, deflationary. It requires no metaphysical baggage. A linguistic philosopher may still talk of three-leggedness being essential to adons, but the point is that the norm that makes it so is not a worldly fact.

4. Conclusion

'*A posteriori* necessities' pose no special problem for linguistic accounts of necessary truth. Neither their *a posteriori* component, nor attempts to show that their necessity has a non-semantic origin, provide compelling grounds for rejecting linguistic accounts. What needs to be captured can be inside a framework that does not presuppose the possibility of substantial necessary truths.

(4) Further Objections to Normative Accounts of Necessity

There are, of course, other objections to the normative account of necessity. Many have been directed towards Wittgenstein's later work, which, I think, contains one of the more convincing accounts. While some of the criticisms have been misplaced and based on erroneous interpretations,[104] others have greater bite and warrant attention. Since I have no intention of defending an interpretation of Wittgenstein, I have no business in responding to critiques of his work that misinterpret him, or that target an account of necessity which is not like mine. But there are many challenges that require a response. Some name Wittgenstein explicitly, while others make arguments that are perhaps designed for accounts inspired by logical empiricism.

This chapter tackles several objections. Four of them are raised by Kalhat,[105] who presents a reasonably fair account of Wittgenstein informed by the relevant literature.[106] Particularly helpful for my purposes, he identifies an account I wish to defend. So, while I may argue in what's to come that some of Kalhat's criticisms misinterpret or misrepresent the account, I mean that they do so in a relatively sophisticated fashion. That is, he has the account correct but he blurs certain distinctions or fails to see its implications clearly. I also tackle other objections that are distinct from those raised by Kalhat, though they often bear certain similarities.

1. The contingency problem[107]

The first objection is a familiar criticism of normative accounts of necessity. Kalhat worries that if we accept (which surely we must) that linguistic norms could be different from what they are, then the necessity that is said to be contained within them is lost. He suggests that for Wittgenstein…

> … the truth-predicate could fail to apply to necessary propositions (in those cases where qua rules they fail to be in force). Yet insofar as these propositions are necessary, the truth-predicate simply could not have failed to apply to them. In calling them 'necessary,'

[104] See Hacker's (2006) rejoinder to Soames (2003), for example.
[105] Kalhat (2008a and 2008b). Kalhat's second paper responds to Glock's (2008) reply to his first. Where my response is largely in agreement with Glock's, I indicate as much.
[106] Baker and Hacker (1985) is, of course, canonical. The final essay of Hacker's (2014) revised edition is of particular importance. 'The Face of Necessity' in Diamond (1991) may also be a helpful corrective to Dummett's (1959) interpretation of Wittgenstein as a 'full-blooded conventionalist'.
[107] I make the same arguments in Couldrick (2022).

> we mean precisely that they could not possibly fail to be true. They could not possibly fail to be true whatever their truth amounts to. (Kalhat 2008a, 9)

Unlike Glock, I do not find this line of argument question-begging.[108] Indeed, I agree with Kalhat that claiming a necessary proposition is one that could not be false is not an expression of realist sentiment but of logic.[109] If Kalhat's argument were correct, it would amount to a direct repudiation of this account. For it holds that while the implementation of specific norms is contingent, the norms themselves, which can be expressed in the form of necessary propositions, nonetheless maintain that rigid necessity. At no point does the contingency at source infect the necessity within. The contention Kalhat takes issue with is the suggestion that saying a given necessary proposition is true amounts to no more than saying the given norm is in force. But is this interpretation correct? He writes:

> … if we took the necessary proposition "All bachelors are unmarried" to be elliptical for the contingent proposition "The rule according to which the word 'bachelor' applies to men who are unmarried, is in force," we would then allow for a possibility that is otherwise (rightly) blocked. We would allow for the possibility that while Kant is a bachelor, we cannot infer that Kant is unmarried because the proposition "The rule according to which the word 'bachelor' applies to men who are unmarried, is in force" is false. (Kalhat 2008a, 8)

Kalhat's argument ignores a central feature of the linguistic account, namely the constitutive role norms play in word-meaning. The possibility that Kalhat claims is allowed by the linguistic account is clearly precluded by it. If Kant is a bachelor, he is unmarried. That is, if Kant is a bachelor, where 'bachelor' has its customary meaning, then he is unmarried. Given the customary meaning of 'bachelor', this is a conceptual truth and it could not be otherwise. One cannot be a bachelor (as we understand the term) and be married. The possibility that is allowed for by the linguistic account is that should 'bachelor' come to mean something different from what it currently does (or what it is assumed to mean in the example), then, of course, one might at some later stage be able to say, truthfully, that 'Pedro is a bachelor and he is married'. And this is clearly the correct result, for we could use the word 'bachelor' to mean anything we like. What the linguistic account does not allow for is the possibility that *our* concept 'bachelor' could be applied to someone who is married, for such an application is precluded by the norms that constitute our concept. This should not be controversial. In cases where it is possible that

[108] Glock (2008, 28).
[109] Kalhat (2008b, 228).

someone is a married bachelor, what counts as a 'bachelor' is different from what we currently count as one. We are talking about different things. Where our bachelors are concerned, there is no possibility that they are married.[110]

Kalhat's confusion can, perhaps, be resolved by drawing his attention to the distinction between words and concepts, where a word is a meaningful sequence of letters while a concept is that which is delineated by word-meaning.[111] The thought behind the first explanation is that it is unnatural to say of a meaningless sequence of letters that they constitute a word. The second is more idiosyncratic, but my terminology makes no material difference to the argument. On this understanding, we would not say there are two words 'bank' (the edge of a river and the financial institution), but that the word 'bank' is used with two distinct meanings.[112] Thus, I want to say that the same word can be used to express two distinct concepts. Why does this distinction help resolve Kalhat's complaint? Well, it is, of course, possible that a rule might not be in force, and in that case the word 'bachelor' might mean something other than 'unmarried man'. But where rules other than ours are in force, it follows that the concept they use 'bachelor' to express is different from ours. Our rules constitute our concept and so where they are not in force, it follows that our concept is not in use. Thus, necessity survives the contingent basis upon which rules come into force via the constitutive role that rules play. Even on the normative account, there remains a sense in which 'all bachelors are unmarried men' is true in all circumstances. For that proposition cannot be false given the concepts employed in it. One might say that while the statement might be counted false in another language, the proposition (what the sentence expresses) cannot be.[113] 'All bachelors are unmarried men' cannot fail to be true so long as those words retain their current meanings – or, in other words, for as long as we continue to talk about the same things with these words. The concept we have now will always be constituted by that norm, regardless of whether we continue to employ it. All the normative account maintains is that we might have used different concepts, meaning different linguistic norms would be in force. Moreover, the statement's necessity is not compromised when the meanings of the words

---

[110] Sidelle (2009, 230-1) provides a similar response.

[111] This is effectively the same analysis given in Schroeder (2009, 93-95), though he prefers to speak in terms of 'sequences of letters' and 'words'. Of course, I do not mean 'sequences of letters' when I say 'words', but the argument is the same. The difference is that when he presents the argument he contrasts a meaningless sequence of letters to the word in question, whereas my comparison case is the same sequence of letters with a different meaning (rather than none).

[112] The counter to this suggestion is that we can define homonyms as two or more words spelt the same with different meanings. My considered view is that our language is flexible on this point. Sometimes 'word' is used as I suggest, other times they can be differentiated by their meaning.

[113] This point is made by Sober (2000, 247), albeit in criticism of conventionalism. He suggests it is also made by Frege (1884 [1968]) as well as countless others, e.g., Ewing (1940), Kneale (1947), Pap (1958), Harman (1967), Lewis (1969), Boghossian (1996).

in that statement change, though we would need a different way of expressing that which is now expressed if those meanings were to change. For what accounts for that necessity is the internal relation between the relevant concepts, and that relation obtains regardless of whether the concepts are in use (i.e., irrespective of whether the norm happens to be in force). An abandoned practice still has rules that constitute it, even if no one plays by them.[114]

A related objection, which can be dealt with in a similar fashion, has been made by Quassim Cassam. He suggests that Wittgenstein's 'weak naturalism' cannot account for the hardness of necessity. Cassam thinks a naturalist with respect to necessity proposes…

> … an account of necessity which makes no appeal to anything outside nature, but insists that his position does not amount to a form of modal scepticism, for a naturalistic conception of necessity is quite compatible with the hardness of the conceptual 'must'. (Cassam 1986, 446)

Cassam's complaint against Wittgenstein's understanding of necessity is that it doesn't allow for the truth of counterfactuals of the kind 'even if everyone believed that two plus two equalled five, it would still equal four'. He thinks that any account of necessity must, as a bare minimum, be capable of counting that type of counterfactual as true.[115] The logic of Cassam's objection is that if our actions are responsible for the norms that are in place, then when our actions change, so must the truth of those norms. Thus, any attempt to reduce necessary propositions to norms of representation fails because those norms cannot hold the line against changes in our behaviour. '2 + 2 = 4' is not a necessary proposition after all, for all it would take for it to be false is for everyone to disbelieve it. This objection is provoked by the following passage in Wittgenstein.

> 'But mathematical truth is independent of whether human beings believe it or not!' – Certainly, the propositions 'Human beings believe that twice two is four' and 'Twice two

[114] It should be noted that Kalhat's dialectic doesn't suggest the argument I have tackled here is the main criticism he wishes to make. Instead, he hopes to build on a criticism made first by Waismann (1965, 66–67, 142), who claims rules cannot be true or false. This is obviously incorrect, as Kalhat acknowledges, for we can say it is true that a bishop may move only diagonally. But Kalhat thinks there is a related problem that cannot be so easily dealt with. Necessary propositions can be combined in conditionals with factually true propositions, but the combination of normative and factual truths creates what von Wright termed 'logical monsters' (1983, 103–209; 1993, 103–113). Thus, necessary propositions are not best understood as norms, for that cannot account for the logical relations between necessary propositions and factual truths without making room for logical monsters. My reason for paying scant attention to this complaint is that I agree with Glock (2008, 27) that 'it remains to be shown that this fear is more than superstition. Why should one not preserve the distinction between factual and normative propositions, while recognising that the common applicability of "true" signals, among other things, that the two can be combined[?]' Statements of necessary truths can be explained in the following terms: 'necessarily, *x*' is true when *x* is required to hold given our semantic norms (cf. Thomasson (2020a; 2021).

[115] Cassam (1986, 448).

> is four' do not mean the same. The latter is a mathematical proposition; the other, if it makes sense at all, may perhaps mean: human beings have arrived at the mathematical proposition. The two propositions have entirely different uses. – But what would this mean: 'Even though everybody believed that twice two was five it would still be four'? – For what would it be like for everybody to believe that? – Well, I could imagine, for instance, that people had a different calculus, or a technique which we should not call 'calculating'. But would it be wrong?[116] (PPF, §348)

It can seem here as though Wittgenstein is conceding too much. For he does not draw our attention to the fact that there is still a sense in which one could maintain that '2 + 2 = 4' is true, namely that the calculus we currently use would still class that proposition as true and it would still be true of that practice. That practice does not cease to exist even if we stop engaging in it. Indeed, '2 + 2 = 4' has a role in constituting it. It would not be that practice if the proposition were false. If the concepts in the proposition '2 + 2 = 4' are those of our current calculus, then there is no possibility of the proposition being false. None of this would be news to Wittgenstein. The explanation for his suggestion that it is possible that people who affirm '2 + 2 = 5' have a different calculus is precisely that our practice is constituted by norms that reject that proposition. Albeit implicitly, the passage quoted does therefore contain within it an answer to Cassam's objection.

In the passage, Wittgenstein is trying to draw our attention to a different aspect, or implication, of his understanding of necessity. Moreover, he is justifiably interrogating the circumstances in which one would assert that counterfactual. For if everyone did believe that '2 + 2 = 5', it seems likely they would be employing a different 'calculus', not making a rudimentary mistake. Indeed, it is doubtful whether they could be said to be sharing our concepts at a time when they affirmed '2 + 2 = 5'. If a community does not correct people when they say '2 + 2 = 5', then the proposition '2 + 2 = 4' is not one of their mathematical propositions and not constitutive of their 'calculus'. So, those people are not making a mistake in saying '2 + 2 = 5' but engaging in a practice different from our own.[117] One interpretation of Wittgenstein's challenge to the counterfactual is to see it as his way of illustrating that there is limited sense in the thought that everyone could follow a rule incorrectly. Their 'calculations' might look like

[116] This passage has provoked readers more sympathetic to Wittgenstein's philosophy than Cassam. Moore, for example, suggests that 'the question about how far our mathematics depends on us […] is one that [Wittgenstein] himself grapples with, very uncomfortably' (Moore 2007, 194).

[117] To what extent we would be willing to call their way of doing things 'arithmetic' or 'calculating', or their practice a 'calculus', would depend on what the rest of their practice looked like. These kinds of questions are considered in detail by Forster (2004). This issue is also considered in Schroeder (2015).

mistakes to us, relative to our standards, but if they are unconvinced by what we say then we might have to accept that they are not engaged in the same activity.

An example from Avrum Stroll might help further demonstrate the point. Suppose a scholar discovers the first book that detailed the rules of chess. Further suppose that the book contained a misprint that made one rule's meaning change and that this misprint has been replicated in every rulebook since.[118] Would it follow that everyone had been playing chess incorrectly? Of course, if what you mean by 'incorrectly' is 'different from the intended rules of the author of the first rulebook', then the answer is 'yes'. But is that what we would call 'making a mistake'? Mistakes are typically made within practices, according to the standards of those practices. Chess, as we play it now, has its own standards that have been adhered to for centuries. The misprinted rule is constitutive of our practice and moves made in accordance with the rule the author intended their game to be played by would count as mistakes in ours. The misprint does not undermine our abilities to apply our own standards. And the rule that was intended by the author of the rules of chess, which was subsequently misprinted, merely defines a distinct set of standards by which one could play a (different) game. We play our game of chess correctly, i.e. in accordance with the rules that constitute it. That we could have played a game with different rules is no cause for thinking we are making a mistake by playing ours. It is absurd to judge moves in one game by the standards of another. After all, I can pick up the ball in rugby without players surrounding the referee and screaming 'handball!'.

Perhaps there are circumstances in which we might be inclined to affirm Cassam's counterfactual. We can imagine, albeit fantastically, a mind-altering drug that temporarily causes everyone to claim twice two is five, but which will soon wear off. Such unusual circumstances may help us understand why someone would want to affirm Cassam's counterfactual, but we should at least acknowledge just how unusual the circumstances are, and would have to be, for everyone to affirm that twice two is five without them having a practice different from ours. And even in such strange circumstances, we can surely still ask whether we can be said to *believe* that twice two is five. Can someone truly be said to grasp the concepts in question if they make such a rudimentary mistake? Without grasping those concepts, we could not ascribe to them such a belief. Can I believe that Julius Caesar was, before he died, a prime number? There is a point in our practices where we will be unwilling to accept that certain things can be believed, and perhaps the question Wittgenstein is posing in the above passage is whether we may reach that point with '2 + 2 = 5'. Of course, we can make mistakes in our calculations, but can we make

[118] Stroll (1994, 112).

such rudimentary ones as these without losing grasp of our mathematical practice? That, it seems to me, remains an open question.

2. The argument from worldly fact

The argument from worldly fact claims that necessary truths are not conventions or facts about conventions and do not, therefore, obtain purely in virtue of facts about convention.[119] Instead, they are made true by the world: 'all vixens are foxes' is true in virtue of all vixens' being foxes. That this statement says all vixens are foxes and the fact that all vixens are foxes provides a full explanation of the truth of that statement.

As Jared Warren suggests, this type of argument has widespread appeal.[120] But it lacks bite without a defence of its premises. Brett Topey, for example, claims it is just obvious that 'all vixens are foxes' is about vixens.[121] Such paucity of argument is difficult to respond to. A charitable reading would be that it is obvious because it is an object-level statement, which at least gives the impression that it is a statement of worldly fact. But to dismiss non-'obvious' readings of 'all vixens are foxes' based on what Wittgenstein calls 'surface grammar' is too fast. As though a philosopher who hears from another room, 'this tele has a mind of its own!', should conclude that its occupant is a panpsychist. And in response to someone who looks at their watch as the night draws in and says, 'is that the time!', should offer to compare watches. But when security tell us we cannot stand here, they are not expressing their scepticism with respect to what we are doing. They are telling us to move out the way. We know this because we understand how they are using their expression. They would accept our transforming their 'you cannot stand there' into 'you are not permitted to stand there' but not 'you are not really standing there'. If we are to be convinced that the factual reading of 'all vixens are foxes' is mandatory, we must rule out its alternatives.

The account I have proposed takes statements like 'all vixens are foxes' to be expressing or reflecting linguistic norms. This is not as counterintuitive a proposal as its critics might

---

[119] I adapt Topey's (2019) rendering of it. His solution to it involves accepting that 'our conventions have the power to make the world one way rather than another' (*ibid*, 1750). There is not space here to explain my disagreement. But one worry is he either commits us to thinking that *a priori* truths, including those he thinks are worldly facts, couldn't have been true prior to the creation of our conventions, or that the creation of our conventions invents facts about the past.

[120] Warren (2015a) suggests Boghossian (1996), Lewy (1976), Pap (1958), Sider (2003, 2011), and Yablo (1992). To that I would add Sober (2000), Williamson (2007) and Hale (2013, 120), who says 'conventions merely determine what proposition is expressed by a sentence – whether that proposition is true is always a further question whose answer is never settled by our linguistic conventions.'

[121] Topey (2019, 1728).

suggest. If asked how we know that all vixens are foxes, it would not be unusual to respond by saying 'because "vixen" means "female fox"'. Indeed, this strikes me as a perfectly good answer. That these norms (or rules) are often expressed in the indicative rather than imperative mood is no obstacle. It is common for rules to be presented in this way: 'white moves first', 'each player is dealt seven cards', 'all players collect $200 for passing Go'. An argument remains to be given to convince us that this interpretation is inappropriate here.

3. What are necessary propositions about?

Kalhat presents a form of the above objection. The thought goes: if necessary propositions are grammatical, or conceptual, then doesn't that make necessary propositions *about* the concepts or words, rather than the objects? Yet, in saying 'a bachelor is an unmarried man', I do not mean that the word 'bachelor' is an unmarried man. Likewise, if I say that blue is a colour, I do not mean that the concept 'blue' is a colour, for concepts do not come in colours. Isn't it, then, absurd to think of necessary propositions as expressing linguistic norms, for such propositions cannot be said to be *about* linguistic items, be they words or concepts?[122] Williamson's gloss on this is that linguistic philosophers take philosophical questions to be 'implicitly about language or thought'.[123]

Now, 'aboutness' here is a somewhat slippery notion. We might do better to stick to the language of constitutive norms. For example, Glock points out that one can have norms that constitute something without the statement of those norms being 'about' that which they constitute. The legal codes that define a country's constitution are not naturally described as being 'about' that constitution.[124] Rather, they lay down norms of behaviour in an analogous fashion to linguistic norms.[125] They state laws that determine what it is for something to be 'constitutional'. However, this case seems to be as much a part-whole problem as something more fundamental. We wouldn't say 'bishops move diagonally' is about chess, but we might say it is about the bishop (in chess). Likewise, conceptual propositions might be said to be about the use of (specific) words.

But even if we grant 'aboutness' to linguistic philosophy's critics, it is still far from clear that much can be made of their alternative vision. Can we make sense of the idea that necessary

[122] Cf. Sober (2000, 254) and Topey (2019).
[123] Williamson (2007, 21). This is how Quine (1960 [2013], 249-50) characterises Carnap's philosophy too.
[124] Glock (2008, 33).
[125] While necessary truths articulate relations that exist given those norms.

propositions are not 'about' words or concepts? Schroeder, echoing an insight that lies at the heart of Wittgenstein's distinction between saying and showing in the *Tractatus*, argues that there really is nothing a conceptual proposition tells us about the objects its constituent concepts are instantiated by.[126] Insofar as 'cygnet' means 'young swan', 'a cygnet is a young swan' tells us no more than 'a young swan is a young swan'. Likewise, one cannot understand the meaning of 'blue' without knowing that it is a colour. So, 'blue is a colour' does not really tell us anything about the colour, for an understanding of that sentence presupposes a knowledge of what 'blue' means (and, therefore, the knowledge that blue is a colour). The underlying thought to both is the following. If we define a 'haldom' as an object that is both red and round, then it is misleading to suggest that information about haldoms is relayed by saying 'haldoms are red and round'. The empirical content of such a statement is nil because the presence of those properties is presupposed whenever something is counted as a 'haldom'. We are already committed to them when we apply the word. The use of 'haldom' carries those properties with it as they are conditions for its correct application.

To those who believe 'self-identity' is a genuine property of all things, this argument may seem to lack bite. But a belief in such a property provides no simple way out of Schroeder's conclusion. If one thinks 'self-identity' is a property of all things, then it is not in the least bit informative to make statements of the kind we are considering. There is no need to reiterate objects' 'self-identity' and doing so would not tell us anything about them (for no object can be distinguished from any other based on a property that all objects possess). The trouble, then, is that appeals to 'self-identity' cannot account for the fact that these statements can be both informative and useful, when used to explain the meanings of words, for example.[127] That necessary propositions might be best explained in a normative fashion therefore survives.

Furthermore, 'self-identity' seems a dubious candidate for a genuine property. When I identify Paolo in the street, nothing is added by asserting his 'self-identity'. The referring term has done the work. He has already been identified as Paolo. It would make no sense to deny Paolo's self-identity and it isn't clear that we are drawing any kind of distinction or describing him as being a certain way. At most, 'Paolo is Paolo' (if not used in an idiomatic sense) seems to be an instance of the principle that explains 'is' as in '=', which in natural language is generally used to indicate shared referents.

[126] Schroeder (2009, 96).
[127] *Ibid.*

3.1 The contingency problem (again)

Kalhat acknowledges that some might be willing to accept necessary propositions are about words or concepts. He does, however, have a separate objection for them. Fortunately, it is a re-run of the contingency problem.

> The trouble now is that metalinguistic statements are contingent. ['All mares are female horses'] says that the word "mare" applies to all and only female horses, but of course, that word could have applied to Californian pineapples instead. (Kalhat 2008a, 11)

As we have seen, this contingency amounts to no more than the possibility that a given community may not share our language. I can say (truthfully) that mares are female horses, and the basis for this is that I understand English. The English word 'mare' means the same as 'female horse'. This is, of course, contingent. But the notion of necessity, as it pertains to 'mares are female horses', never concerned this aspect of the norm. Rather, the necessity arises from the constitutive relation between a linguistic norm and the concepts it (partly) constitutes. It will always be true that, according to the concepts we have now, all mares are female horses, even if those words later come to express something else, or we cease to employ these concepts altogether.

Kalhat considers this response briefly but rejects it out of hand as he thinks that, far from constituting the concepts, these statements now presuppose them.[128] But statements of rules were never supposed to constitute concepts. It is the rule itself that does. A concept is defined by its semantic rules, not by the statement of those rules.[129] It is the way a word is used that determines meaning, not a report – or statement – of that use. Necessary truths are unassailable because our use of the relevant concepts requires us to accept them. And for the question to arise the concept must be used, presupposing the norms that constitute it. Furthermore, to say something 'will always be true' (or words to that effect) is only another way of saying: this is the shape of the concept we are talking about. A different shape would make for a different concept, even if the underlying similarities were such that we would be inclined to count only one concept (something can be altered without its identity changing).[130] At the very least, however, it would not be *this* incarnation (or variant) of the concept – something, after all, would have changed. Insofar as we are seeking to identify the specific incarnation of the concept

[128] Kalhat (2008a, 12).
[129] See Baker and Hacker (2014, 46-55) for a detailed discussion of the differences between rules and their formulations.
[130] This is merely a reflection of the diachronic identity of concepts. In most cases, I suspect, our concepts can tolerate some alterations before ceasing to be those concepts. This point is discussed in Schroeder (2009, 94).

we have *now*, the norms that presently constitute it cannot fail to be in place. They hold for this incarnation come what may because they are what make it the concept that it is.

We might think of this as there being internal and external aspects of conceptual propositions. The internal aspect is that which has the character of necessity: the constitutive relation that obtains between a norm and the concepts it (partly) constitutes. The external aspect concerns whether the norm is in force and allows us to recognise that there are alternatives.[131] Schroeder introduces a similar distinction between the internal and external perspectives of a given practice.[132] From within the practice, there is no alternative to the rules that are followed, for to be engaged in that practice is to follow the rules that constitute it. The external perspective, however, acknowledges that we may choose to no longer follow those rules and simply do something else.[133] (One must only move the King one square at a time, but that doesn't stop one from throwing it across the room if so inclined.)

3.2 'Just words'

The original objection of this section is sometimes expressed as a kind of disappointment from philosophers who want more than what the linguistic philosopher is (allegedly) prepared to offer.[134] So, David Papineau suggests that the philosophical analysis of concepts 'may tell us that *if* there is a propositional attitude that requires truth, justification, and so on, then it is knowledge', but laments that 'this seems far less than we actually get from the relevant thought-experiments.'[135] He takes…

> Gettier to have shown not just that our concept of knowledge imposes a requirement of non-accidentality, but far more interestingly that this requirement is satisfied by real knowledge – that is, the state that plays an important role in the world and is displayed in many paradigm cases. Similarly, I take Kripke to have shown not just that we conceptualize names causally, but in addition that real name-bearer pairs - all those many instances we are familiar with - are causally related. (Papineau 2009, 19)

---

131 Carnap invokes a similar distinction between internal and external perspectives. See Shapiro (2000, 342-3).
132 Schroeder (2009, 95).
133 Coffa (1991, 139) makes essentially the same point when he says 'convention, semantically interpreted, is merely the opposite side of necessity. In the range of meanings, what appears conventional from the outside is what appears necessary from the inside.'
134 Or that 'the old, confining orthodoxy of philosophy as linguistic analysis', as Soames calls it (2006, 307), has to offer.
135 Papineau (2009, 19).

Papineau thinks our interest should be in the state (knowledge) that plays an important role in the world and is manifest in paradigm cases. But where does the impression that linguistic philosophers are not (or cannot be) interested in just the same thing come from? When we investigate how our language works, we are not studying an abstract logic independent of the world's happenings. On the contrary, we are investigating a language that is embedded in our lives. We study what we do with language, what we apply our words to, and when we are willing to count an object or state of affair as instantiating a certain concept or description. We investigate the conceptual framework of phenomena that play important roles in our world, and this means paying particular attention to paradigm cases.[136] J. L. Austin probably put the point best by saying:

> When we examine what we should say when, . . . we are looking again not merely at words . . . but also at the realities we use the words to talk about. (Austin 1957, 8)

Papineau claims that we are not interested in the question of whether we 'conceptualize' names causally but whether the relation between a bearer and its name is causal, just as we are interested in that important state *knowledge* not our concept 'knowledge'. But the linguistic philosopher interrogates this distinction. There is obviously a difference between a concept and that which instantiates it.[137] But that 'knowledge' tracks some hidden entity, the essence of which we wait to discover, is what they deny. We investigate *knowledge* by studying cases that we would count (or not) as 'knowledge'. And to do so we rely on our semantic expertise. That is why often philosophers appeal to what we would say in different circumstances.

> It is true that many philosophers […] make a point of rejecting the appeal to 'what we say', but this rejection is often undermined by their way of discussing the matter. In such discussions we are commonly invited to agree that in such and such a case a person would or would not know that p (that 'it is clear' he would not know, etc.); but the only way to assess such claims is by considering whether a person would or would not *be* said to know that p in those circumstances. (Hanfling 2000, 96)

Does this mean we cannot be wrong about what knowledge or naming involves? Of course not. The goal of linguistic philosophy is not to discover how we 'conceptualize' names, if that means how we pre-reflectively think they function. For without thorough investigation, our assumptions regarding how names work may turn out to be wrong. One can have a certain picture of how something is, what the nature of a given object is, only to realise from a

[136] See chapter nine for more on the embeddedness and importance of concepts.
[137] See Hanfling (2000, 17).

conceptual investigation that this picture is misplaced. To realise that naming, for example, is a much more diverse phenomenon than one had first thought.[138] Indeed, one classic source of philosophical problems arises from disparities between how we think about a given concept and how we actually use it. Hence a linguistic philosopher may be inclined to agree with Papineau that philosophy is not interested merely in how we 'conceptualize' names, but in how they really function. By studying conditions of conceptual application, one investigates what it takes to fall under that concept and so be an object of that kind. As Hacker suggests:

> 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 X-s would have struck analytic philosophers of the [twentieth century] as perverse.[139] (Hacker 2009, 338)

One of the motivations behind the linguistic conception of philosophy is that it is difficult to understand what philosophical investigations into the nature of knowledge and naming could be if not an investigation into the kinds of phenomena we count as falling under those concepts. How would we know what types of cases to consider otherwise? How do we know that Gettier cases don't constitute knowledge? It's not as if knowledge is a substance we can study and ask what it consists of – what its inner constitution is like. Even in the case of substances, as I argued in chapter three, empirical investigations do not reveal what *really* counts as a particular type without an accompanying *a priori* principle. Essence is only to be discovered via empirical study if there is an *a priori* principle that stipulates as much. Whether that principle really applies in each case is a question akin to what the principles of knowledge are. Could something be the same substance despite having a different atomic structure? Does knowledge require belief? The methods by which we may obtain answers to these types of questions carry similar constraints. As Wittgenstein puts it, 'our investigation […] is directed not towards phenomena, but, as one might say, towards the "possibilities" of phenomena'.[140]

The opposition of word-meaning and the nature or essence of something is specious because we cannot identify what's in question independently of word-meaning. We know which phenomenon to study because we understand word-meaning. The attempt to drive a wedge between the nature of *knowledge* and the meaning of 'knowledge' is rendered absurd by the observation that *knowledge* is just that which falls under our concept.

---

[138] This is the lesson I take from Wittgenstein's discussion in the *Investigations*. Relatedly, Raatikainen (2020) has suggested there are at least four different types of referring expressions (and shows that many proponents of the new theory of reference accept that not all referring terms work the same way).
[139] Cf. Hanfling (2000, 17) and PI §370.
[140] PI §90.

This response does not alter the common trajectory of the philosophical dialectic, from folk understandings to greater clarity. But it does reimagine philosophical 'theory'. Philosophical 'theories', on this account, are often attempts to describe conceptual structures (real or imagined). Competing theorists come to define concepts in different ways. This no doubt raises many questions. One is whether the introduction of new concepts through fresh 'theories' does more than avoid previous philosophical problems, rather than, say, solve them. This is not to say there is no place for the proposal of new concepts. Indeed, similar (though different) concepts may help to illuminate existing conceptual practices. But one might wonder whether philosophical problems are really solved by introducing concepts that did not give rise to the problem in the first place. Instead, we might need an accurate characterisation of the practice in which the problem was first encountered.[141]

4. Boghossian and pre-linguistic times

There are two more objections that bear similarities to those already encountered. The first is presented by Paul Boghossian. He criticises what he calls 'metaphysical analyticity'. '[A] statement is [metaphysically] analytic provided that, in some appropriate sense, it owes its truth value completely to its meaning, and not at all to "the facts."'[142] While I have deliberately avoided talk of 'analyticity', the questions he raises are relevant. For they are largely questions regarding whether appeals to language (as opposed to 'the facts') are capable of accounting for, and explaining, necessary truths.

One question is: '[h]ow 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*?'.[143] Linguistic conventions account for what is meant by an expression but seem powerless to make what is thereby expressed true. The point here is the seemingly innocuous observation that a statement is true just in case it says '*p*' and *p*. But this can be made consistent with my account. Insofar as one is licensed to say 'necessarily *x*' just in case '*x*' is a constitutive norm or its logical consequence, then 'necessarily *x*' is true when

[141] Cf. PI, §132. This was famously Strawson's (1963) criticism of Carnap's explication in philosophy. See Pinder (2020) and Preston and Schroeder (2020) for a discussion and defence of Strawson's arguments. Of course, it might be enough to simply avoid problems, rather than resolve them, and there are certainly benefits (in the right circumstances) to developing more precise, clearly-defined concepts.
[142] Boghossian (1996, 363).
[143] *Ibid*, 364.

that condition is met.[144] 'Vixens must be foxes' is indeed true just in case that is what a vixen must be (just in case that is what the norm requires).

Boghossian goes on to make an argument against the notion that conventions generate necessary truths. He wonders whether we are really expected to believe that, prior to a meaning being stipulated for the sentence 'either snow is white or it is not', snow was not either white or not white. While this might be a problem for those that claim our adoption of certain conventions makes the relevant proposition true,[145] it is unclear that it troubles the normative account of necessity. Insofar as the concepts and logic involved in 'either snow is white or it is not' are as the example assumes, we are entitled to say it is always true. The reason for this is not that it represents the *a priori* structure of the world, but because it is a timeless truth. It faithfully represents the logical structure we use to think about the past. When Boghossian suggests it is overwhelmingly obvious that the statement was true before an act of meaning, he is correct. In pre-linguistic times, snow was either white or not white. The important omission is that our thinking about pre-linguistic times relies on conceptual, logical, and mathematical norms. The structure of our thoughts presupposes these principles. While the statement might appear to concern pre-linguistic times, what it represents is a logical truth. It reflects the laws of excluded middle and non-contradiction with respect to the proposition 'snow is white'. One can likewise say that 2 + 2 equalled 4 and nothing could be red and green all over simultaneously in pre-linguistic times. Expressions of these norms can be regarded true at all times precisely because any time those concepts are involved, these norms are presupposed. While there was no one to count the proposition as true in pre-linguistic times, the point is that as soon as the question arises the relevant proposition must be regarded as true.

There may be a lingering air of paradox here. Boghossian stipulates that he is talking about a time in which the relevant language has not been created, while I claim norms that belong to that language nonetheless account for the statement's truth in those pre-linguistic times. But this paradox can be dispelled. Consider the following empirical statement: 'Mont Blanc existed before humans did'. This statement happens to be true. Mont Blanc massif was completed around fifteen million years ago. But does our acknowledgement of this fact eschew the use of today's language altogether? Clearly not. Of course, there were mountains prior to the

[144] A point made by Thomasson (2020a, 86-7) who develops this in more detail. Similarly, Hacker explains that 'to say of an expression of a rule that it is true is simply to *affirm or concede the expression of the rule*, just as ascription of truth to a proposition "*p*" in the utterance "It is true that *p*" is simply to affirm or concede the statement that *p*. In both cases, "It is true that *p*" is equivalent to "*p*"' (Baker and Hacker 2014, 279).

[145] Although advocates of this approach might argue that our current conventions make the proposition about the pre-linguistic past true. Nyseth (2021) hints at the possibility of a similar response.

development of our concept, but what we mean by this is that there existed things that met the standards necessary to be considered a 'mountain' before those standards were devised. We still appeal to those standards when we ask, 'were there mountains prior to the existence of language?'. Similarly, when asked whether snow was either white or not white in pre-linguistic times, the answer is obviously 'yes'. But that answer no less relies on the principles we now use to think about those pre-linguistic times than the question about mountains does. When a statement accurately reflects such principles, it will be regarded as (necessarily) true.

This is little more than an acknowledgement that our language can be applied beyond our presence: we can talk about pre- and post- linguistic times, as well as corners of the universe where we will never set foot. Our presence may be required for the describing, but certainly isn't for the thing described. Likewise, our presence may be required for the development of modes of description, but once developed their application is not limited to times or places we have been or will be present.

5. Nyseth and the truth-contrast thesis

Fredrik Nyseth modifies Boghossian's objection and targets what he calls the 'truth-contrast thesis', the claim that 'necessary truths are fundamentally different from contingent ones since they are not "made true by the (worldly) facts"'.[146] While he appears to anticipate (at least in general terms) my response to Boghossian, he thinks there is another problem in the vicinity. He does not argue against the linguistic explanation of *necessity* (which he is sympathetic to), but only against its account of how these propositions are said to be true. He thinks that necessary truths, though guaranteed to be true by linguistic norms, are made true by the facts.

'Snow is white or snow is not white' is his main example. He claims that those who subscribe to the truth-contrast thesis must provide an additional explanation for how the statement is true. This is problematic, he suggests, because the normal explanation for disjunctive statements' truth, that at least one disjunct is true, seems to account for this case. In fact, given that one of the two disjuncts will always hold, it will always be true according to the normal criteria. Hence, we do not require another explanation for how this statement is necessarily true. Moreover, he argues, the alternative explanations are unable to account for there necessarily being some fact to make the disjunction true.

[146] Nyseth (2021, 264).

An analogous argument to Nyseth's can be made in a case that unquestionably relates to normative rules. Consider a game played by rules including the following:

(i) It is a two-player game.

(ii) The game is only complete when one player wins.

What are we to make of the following statement?

(iii) One person wins the completed game or the other does.

We could, I take it, make the same argument Nyseth does. The statement is true when at least one of its disjuncts is (i.e., if one of the two people wins the game). Given that, if the first person doesn't win the completed game, the other will, one of these disjuncts will be true in all circumstances. The disjunction is made true by this fact.

But is it obvious that this is a worldly fact? It seems to merely reflect the rules of the game. If people are playing the game, it follows trivially that one must win for it to be complete. That is part of what their completing the game consists of. One aspect of their completing the game, of what it means to complete the game, is that one of them is declared the winner. That one will win the game, however, *is* an empirical proposition, for the game may never be completed.

For (iii) to be true, it must say something. This requires that the game and its completion be determined. But once they have been determined by the rules contained in (i) and (ii), (iii) is no more than a restatement (or reflection) of those rules. Indeed, it conveys no information about the world. At most, (iii) expands on what is involved in the playing and completing of that game (i.e., it explains some of the rules). Suppose Harry and Sam complete the game. Again, one of them must have won. But this is not an additional worldly fact. On the contrary, once we understand the game in question, there is no information contained in the proposition beyond the fact that they have completed the game. This has clear parallels with the argument we saw Schroeder make in response to Williamson earlier.

As for the claim that this sort of explanation is incapable of explaining *why* one of the disjuncts must hold, that is clearly not the case. It seems perfectly reasonable to explain that Harry or Sam will win the completed game by appealing to the rules that state one of the two players must win. There is no coherent description of a completed game in which one of them does not. Rules legislate that the game is only complete when there is a winner. We do not count a game as complete if one of them has not won. This norm is logically prior to any statements about completed games because it is a condition of the possibility of such games. Hence, facts

about completed games cannot vindicate the game's constitutive rules: what counts as a 'fact about completed games' is determined by those rules. To suggest otherwise is to put the cart before the horse.

If this is right, the mere fact that it *seems* the normal explanation for disjunctive truth holds isn't sufficient support for Nyseth's argument. Moreover, the claim that accounts like mine cannot explain why one of the disjuncts must be true is incorrect.

6. Conceptual norms and independent existence

A third line of argument Kalhat pursues against Wittgenstein's approach is again related to Boghossian's. It appears to turn on a misunderstanding of the relation between a concept and the object the concept refers to. He claims that:

> Wittgenstein regards ['Nothing can be red and green all over'] as constitutive of the nature of the colours red and green. It follows, therefore, that the nature of the colours red and green is of our own making […] But the colours themselves are surely not of our own making. How, then, can the nature of colours be dependent on us when the colours themselves are not?[147] (Kalhat 2008b, 230)

> … if we have not created the colours red and green, how could we have nevertheless created their nature? For what sense attaches to the idea that red and green could have existed without having the properties that make them the colours that they are? That would be to say that red and green could have existed without being them. (Kalhat 2008a, 13)

To be clear, Wittgenstein did not think that colours were of our own making in Kalhat's sense. That would lead to an absurd form of idealism, which is rightly blocked.[148] We can maintain that the nature of colours is determined by our linguistic norms, while also holding that colours have an existence independent of us. Just as before there were any humans on earth, and so before colour-grammar and calculus had been invented, it was still true that the earth had only one moon and that much of the grass was green. There was just no one around to say it (or introduce

[147] Cf. Canfield (1981, 16).

[148] Of course, Anscombe (1981) famously argued that Wittgenstein was an idealist with respect to certain domains. This has been repeated by Bloor (1996). See Dilman (2002, 57-82) for a convincing rebuttal. A more difficult, and complex, interpretation is presented by Williams (1981) and Lear (1982; 1984), who argue that the later Wittgenstein was a transcendental idealist. Such an idealism is not the kind I refer to here. For a rejection of this interpretation that nonetheless shares some sympathy for it, see Moore (2007; 2012). For a less conciliatory approach, see Mulhall (2009a). Chapter ten deals with this form of idealism explicitly.

the concepts necessary to say it). One mistake Kalhat seems to make, judging from the second quotation, is conflating logical and temporal priority. For something to count as 'red' or 'green' does of course require a colour-grammar that defines those concepts. But it does not follow that anything that existed prior to colour-grammar could not have been red or green. For what it means to be 'red' or 'green' does not depend on language in the same way that 'speaking English' does. Once one has adopted standards with respect to one's colour-concepts, one can begin to explore where those properties have been instantiated.

Overall, I am in broad agreement with Glock's response to Kalhat. As such, I will focus on what Kalhat has to say in reply. But first let us consider Glock's argument.

> … what does it mean to say that colours exist independently of us? It means that the concepts we use do not alter any of the pertinent facts about objects, notably their visual appearance. For instance, they do not render hitherto colourless objects coloured (red, green, etc.). Yet those facts can be stated only in terms of concepts which, at least according to Wittgenstein, are of our own making. There is an indefinite number of facts involving objects, which could in principle be stated; for there is an indefinite number of properties that objects possess. Only some of these properties are such that creatures of our cognitive and perceptual capacities can notice or ascertain them. And of these, only some are captured by our concepts, and hence feature in our statements. (Glock 2008, 30)

That some objects exemplify various combinations of properties is not something grammar dictates. All grammar does is give us a means by which we can identify those properties and use them in our thought and talk. Grammar may determine what it is that makes some given property *that* property, but it is not responsible for that property's existence. Where and when a property exists is wholly independent of the grammar that furnishes us with the tools to think about it. Grammar's role is only to determine what it means to say, 'object X has property Y'.

Kalhat, in his reply, takes issue with the next step in the argument. Glock claims that while there are empirical facts that influence the concepts we have, nothing – in principle – would stop us from employing colour-concepts which do not abide by the norm 'an object cannot be red and green all over at the same time'. Kalhat suggests this is mystifying. We can indeed imagine cases where we might loosely talk of an object being red and green all over at the same time, such as when an object oscillates between red and green from moment to moment. But that object would not be red and green *at the same time*. Glock, so Kalhat claims, 'needs to

show that it is possible to describe an object as simultaneously red and green all over in a strict manner of speaking, i.e., in accordance with the established meaning of the terms involved.'[149]

But it is Kalhat's counterargument that is mystifying. For what he effectively claims is that he cannot conceive of a scenario in which 'an object cannot be red and green all over simultaneously' would be false. But surely the surprise would be if he could describe such a case! Of course, one cannot describe a situation where that proposition would be false, for it is a necessary truth. It is a rule of representation and so counterexamples are, by definition, ruled out. A feature of 'an object cannot be red and green all over simultaneously' being a conceptual proposition is that one cannot describe a situation that conflicts with it. A proposition that appears to deny it must be using some of the words to mean different things. For the sense those words have is partly constituted by the norm expressed in the proposition. Kalhat appears to anticipate an argument of this kind but suggests the 'complaint is misplaced because in the cases just considered it is non-colour words that are not being used with their established meaning (e.g., "oscillation")'.[150] Moreover, we are told that the charge of question-begging cannot be made by Glock since he thinks the meanings of 'red' and 'green' might survive the abandonment of 'nothing can be red and green all over'.[151] This response is misguided.

Firstly, it might be non-colour words doing much of the work in 'an object cannot be red and green all over simultaneously'. For example, it might be that, in the context of colours, we effectively mean 'uniformly coloured' by 'all over', i.e.: not one bit of the object is another colour, although we should acknowledge that the 'all over' claim can be restricted to, say, a particular surface. If after I said of an object that it is green all over you chipped the paint to show a different colour underneath, I might reasonably claim I hadn't meant to say anything about what lay beneath the surface. Our abandonment of 'nothing can be red and green all over' may therefore largely concern the notion that an object can be uniformly coloured. In a world where colours rapidly oscillate, we may have little use for this way of speaking.

I think the origin of Kalhat's mistake is that he imagines the rejection of a conceptual proposition to be something greater than the disuse of some of the concepts contained within it. In this regard, it might be helpful to invoke the distinction that we saw Moore make earlier.

---

[149] Kalhat (2008b, 231).
[150] *Ibid.*
[151] *Ibid,* nn. 1.

> To reject a proposition is to decline to think in such terms: it is to repudiate some or all of the very concepts involved in the proposition. To deny a proposition, by contrast, *is* to think in such terms, but to count the proposition false. (Moore 2012, 593)

Conceptual norms may be rejected, but never denied. The point of the scenarios Glock imagines, such as cases where objects oscillate rapidly between different colours, is to motivate the employment of different concepts, not to imagine a case where our grammar is *falsified*. The grammars we end up with are not arbitrary in the sense of lacking explanation or justification. The world around us, our physical makeup, culture, and interests influence the way we come to make sense of the world. Whatever physical constraints are imposed on us by the world, they do not account for the necessity in propositions such as 'an object cannot be red and green all over', for, in contrast to necessary propositions, those physical constraints could have been otherwise, meaning we can describe states of affairs where they are different from what they are. Were significantly different physical constraints in place, some of our existing concepts might cease to be useful. But that would not prevent us from imagining circumstances where such concepts would be useful (and used).

The second criticism Kalhat makes in his reply is aimed at the distinction Glock draws between the 'the causal properties which are responsible for the way coloured objects affect our perceptual apparatus' and the essence, or nature, of colours as expressed in grammar, 'which give rise to *a priori, conceptual inferences*.'[152] This distinction maps onto the one I have drawn above, between features that influence the kind of grammar we have and the necessity found in conceptual propositions. The linguistic account need not deny the role causal properties have in shaping our grammar. It need only deny that such causal properties are fit to account for the necessity in norms such as 'an object cannot be red and green all over at the same time'. Kalhat claims that the physics of the world dictates that an object could not be red and green all over at the same time, but for that to be taken seriously he would need to adequately describe the case that physics allegedly rules out (something we have already seen him rightly suggest is impossible). Without an adequate description of such a case, an object's not being able to be red and green all over cannot be understood as a limitation, or constraint, placed on it by the laws of physics, for there is nothing that object cannot be and therefore no constraint in place.

7. Unaccounted-for modal connections

[152] Glock (2008, 30-31).

The final objection Kalhat raises against Wittgenstein's approach concerns the kind of explanation provided by a linguistic account. Insofar as the linguistic account seeks to reduce necessity to normativity, he thinks it fails. The normative explanation relies on the modal notion it seeks to explain. As Glock points out, the potential problem for these accounts is not so much a frustration of their proponents' reductionist desires (often they have none), but rather the implication that some modal notion is left unaccounted for 'and which perhaps cannot be accounted for in linguistic terms.'[153] Modal notions need not be reduced to more basic non-modal terms, but a normative account of them must be forthcoming. Kalhat's argument doesn't, I think, undermine this possibility.

> To follow the rule correctly is to apply it in a way that is compatible or consistent with the rule, and to fail to follow the rule correctly is to apply it in a way that is incompatible or inconsistent with the rule. […] If all possible linguistic behaviour were consistent with a rule, then it would just not be a rule. Therefore, in order for something to be a rule, some of its applications must be consistent with it, and some must be inconsistent with it. But the notion of consistency is evidently a modal notion. (Kalhat 2008a, 20-21)

Kalhat's conclusion does not follow from its premises. He rightly says some linguistic behaviour must be inconsistent with a rule for that rule to count as such but concludes that for something to be a rule some of its applications must be inconsistent with that same rule. Whereas all that really follows is that some *possible* applications or behaviour must be inconsistent with the rule (not that they are, in fact, applications of the rule in question). And this takes us to the crux of the issue.

The correct response is one that Kalhat considers but (wrongly) dismisses. His objection suggests that a rule precedes its application. We first have the rule, and we can then apply it in different ways. Wittgenstein rejects this picture. 'To follow a rule' is a success-verb.[154] Whether someone is following a rule is determined, in part, by how they apply it. The relation between a rule and its application is internal. Rule-following is a technique and whether you are employing the technique will depend on how you act. A rule stipulates what counts as following it. For instance, if the rule is that bishops move only diagonally, then it is only by refusing to move the bishop non-diagonally that one can follow the rule. Likewise, if there are conditions under which something counts as 'round', then objects that meet those conditions are round (and those that don't are not). Where conditions are unclear, it may well be difficult to determine whether

[153] *Ibid*, 39.
[154] Baker and Hacker (2014, 137). See 135-40 for a wider discussion of 'following a rule'.

something meets them. But rather than that implying there is a gap between the conditions and their application, it instead speaks to the very fact that the conditions are themselves vague (hence it is not obvious what does and doesn't count as following the relevant rule). There is, in other words, no gap between a rule and its application. And there is no residual necessity left unaccounted for.

Kalhat suggests in response that Wittgenstein is wrong about the relation between a rule and its application being internal.

> … this reply cannot be right for it appears to entail that there is no such thing as following a rule incorrectly. We must surely allow for the possibility of the child who is learning how to apply the rule "+2," and in the process of doing so makes all manner of mistakes (e.g. at one point produces the series "0, 2, 4, 9, 14, 57 . . ."). He is trying to follow the rule "+2" correctly, but fails to do so. He applies it in a way that is inconsistent with the rule. (*Ibid*, 21)

Nothing in what I have said suggests one cannot try, and fail, to follow a rule. That someone is trying to follow a rule is determined by their intentions. In the case of the child, the context makes it clear that they are attempting to follow our rule. The point is they try – and fail. One can describe this either as following the rule incorrectly or just failing to follow the rule. There isn't some liminal state between following a rule and doing so correctly or incorrectly. We don't first follow the rule, then follow it correctly or incorrectly. Rather, we try to follow the rule and sometimes succeed, sometimes fail. One is not following the rule and doing so incorrectly when one gets things wrong but trying to follow the rule and failing to do so.

An analogous case would be one in which I instruct you to turn left at the lights but instead you turn left at the garage. In that case, you haven't followed my instructions. We might say you failed to follow them correctly, if you intended to follow them but failed to (because you misremembered or misheard them, etc.). The key point, however, is that it makes no sense to say: 'I followed the instructions but didn't do as they say'. Of course, you can *think* you've followed the instructions. Just as you can read a recipe line by line and think you followed it to the letter – only to realise you haven't because you, for example, missed an ingredient. Upon realising this, it wouldn't make sense to claim you had still followed the recipe even when failing to do as it says. Kalhat's argument relies on there being some notion of following a rule independently of whether we do so correctly or not. But trying to do something is not the same as doing it. We can distinguish between trying to follow a rule, thinking we are following a rule, and actually following it.

The incoherence of Kalhat's argument is laid bare in the final sentence of the quoted passage: 'He applies it [the rule] in a way that is inconsistent with the rule.' At most this can mean: he has failed to apply the rule, for what he has done is inconsistent with it. This might seem to re-open the door to Kalhat's objection for it appeals to a modal notion, namely inconsistency. But there is no problem with using modal notions in this context, so long as they can be accounted for normatively. And something's being inconsistent with the rule is just another way of saying that one is not following the rule when performing that action.

This takes us to a point I will expand upon in the following chapters, namely that the constitutive role norms (or rules) play is what accounts for necessity. When one says that vixens *must* be female, for example, that is tantamount to saying something wouldn't be a vixen if it weren't female. The reason for this is the norm that determines what does and doesn't count as a 'vixen'. This norm can be made explicit by using modal vocabulary: 'a vixen can be… but it can't be…'. The modal license is provided by the norms that determine what it is to be a vixen. It is not necessary to use modal language, however. Wittgenstein makes this much clear in *Zettel*.

> Do not say "one cannot," but say instead: "it doesn't exist in this game". Not: "one can't castle in draughts" but – "there is no castling in draughts" (Z §134)

Kalhat is dissatisfied by this for reasons that will become clear. Wittgenstein's point is one I made in the first chapter. Saying 'one cannot' makes it sound as if there is *something* that one cannot do. But the point is that 'castling in draughts' doesn't represent anything. It isn't a move in the game. It is an empty description. The 'cannot', then, is more akin to not being able to see something because it isn't there, rather than it being a limitation of ours. 'It doesn't exist in this game' and 'one cannot do that in this game' are really two sides of the same coin, but the former guards against our conceiving the situation (wrongly) as exposing a limitation of ours.

Kalhat, however, argues that these explanations are not equivalent. The modal explanations are more basic:

> … the reason why "The Tower is red and green all over" does not exist in the language-game is that it is impossible to utter the sentence, and use the words "red" and "green" with their customary meanings in that game – not the other way around. It is the impossibility of the move that explains why "we do not allow" it "in our linguistic practice", and hence why the move does not exist in the game. (Kalhat 2008b, 234)

This view has significant flaws. Disallowing certain moves is a way of drawing boundaries around concepts and so determining meaning. Yet Kalhat claims that it is the impossibility of

combining certain meanings which is primary. The problem is that there would be no meanings to speak of were it not for our having allowed or disallowed certain moves – without us having decided which moves would exist in this game. The impossible moves he discusses are already accounted for in the norms that constitute the relevant meanings. They are not external to them. He continues:

> … the reason why you have to perform a certain move (refrain from making a certain move) in a game is that making the move (refraining from doing so) is *necessitated* by a rule that is constitutive of the game: so long as you want to play the game, you have to make the move (refrain from doing so). […] to say that a move goes against the rules is to say that the move is inconsistent with the rules. The move is inconsistent with the rules, and the rules are constitutive of the game; hence, it is impossible to play the game and make that move. It is this residue of modality that Wittgenstein's account leaves unexplained. (*Ibid*, 235)

Kalhat thinks that the fact we cannot make a certain move is distinct from the constitutive rules of the game. Nevertheless, the rules do *necessitate* that fact. Hence, the residue of modality is located between the rules and the moves one can and cannot make as a result. The thought is that some necessity is left over *between* a rule and its application, which cannot therefore be accounted for by that rule. His analysis rests on the spurious claim that applications of a rule are independent of the rule itself.

When he claims that the reason an object cannot be red and green all over is that it is impossible to utter the sentence and use the words with their customary meanings, one ought to reply that using words in that way would amount to using them without their customary meanings. Linguistic norms express regularities within our language. To use words outside of their established regularities just *is* to use words without their customary meanings, for those established regularities are constitutive of that meaning. If we take seriously the idea that norms are constitutive of word-meaning, we have an explanation for the impossibility, namely that using words in an unusual way is to use them to mean something else – or nothing at all (for their meaning is constituted by the norms of their established use, norms which are flouted when words are used in unusual ways).[155]

[155] This point can be overstated and perhaps I have done so here. We often project words into new contexts without failing to make sense. But such projections, insofar as they are new, do breathe a lease of life into our words that was previously missing (I use metaphor here advisedly). Their meanings are thus extended (or enriched). Arguably what is more interesting is the fact that it is often difficult, sometimes perhaps impossible, to determine whether some particular projection is a novel use of a word or not – whether we are extending the use of a word or simply

Kalhat thinks the debate boils down to which explanation is the more basic. The main point to be made, however, is that the impossibility of combining 'red all over' and 'green all over' can be explained by the constitutive relationship that obtains between linguistic norms and meaning. The basic explanation is to be found in the constitutive role norms play, which entails that using words outside of their usual bounds breaks those words from their established meanings. The reason it is impossible to breathe oxygen without breathing in an element with eight protons in the nuclei of its atoms is that oxygen *just is* that element. The basic explanation is the constitutive relation that holds between the two, which can be (and often is) articulated in modal terms.

8. Quine's regress and logical consequence

It is worth ending with an objection to this account of necessity from the realm of logic. While logic is not my chief concern in this project, the response to this objection clarifies the role statements of norms have (and don't have) on this account. Moreover, it is possible for someone to think that some statements are – in a sense – true by definition but hold that they nevertheless represent fundamental logical truths (and those truths, they might think, are substantial). It is therefore worth considering what has been taken to be a knockdown argument against 'conventionalist' accounts of logic.

In 'Truth by Convention', Quine presents an argument against conventionalism that might be thought to apply to my account. He presents a dilemma: either one thinks we must stipulate every logical truth, which leaves an impossible task for finite creatures, namely stipulating an infinity of logical truths; or one thinks we stipulate several *general* logical truths that can then be applied to infinite cases. While the latter may seem appealing, it is scuppered by the fact we require a further logical inference to move from a general truth to a specific one.[156]

Quine believes this argument only rules out views on which logical rules must be explicitly stipulated. One reason for thinking this is that if rules of inference are implicit in our behaviour, then what we might have learned is a technique for transforming propositions that

applying the word with its customary meaning in new circumstances. This subject will be tackled in Part II. I should note here, however, that even if in some cases it is indeterminate as to whether some use of a word is novel, it does not counter my argument. For here I am concerned with norms that are quite determinate and to which we adhere strictly. Where rules seem to grant several exceptions, the matter is of course complicated. But in such cases, it is not clear that one has flouted a norm governing the use of the word, hence the difficulty in determining whether to count the case as new or not. The mere existence of these cases does not, therefore, spell trouble for my argument.

[156] Quine recognises a similar regress is found in Carroll (1895).

can be applied to an infinity of cases. The inference from general rules is not required in the act of making particular inferences, though we might explain the legitimacy of those particular inferences by showing that they follow from the general rule. Although Quine spends some time building the dilemma for explicit conventionalism, he presents a much simpler reason that shows that explicit conventionalism is hopeless, when he suggests that the communication of the relevant logical truths 'depends upon free use of [the] very idioms which we are attempting to circumscribe'.[157] Explicit stipulation is a non-starter for this reason. No language could begin with explicit stipulations regarding its rules because there is no language in which those rules could be stated. This is, as Richard Creath says, 'little more than an instance of the more general truth that without a language you cannot say much of anything.'[158]

Quine is also sceptical of relying on inexplicit rules, albeit for different reasons. In short, without the act of deliberately adopting rules, he sees little reason to regard these inexplicit rules as anything other than firmly held statements. This objection is unconvincing.[159] As Warren suggests, 'speakers react in a uniform way to cases that they not only have not encountered, but could not have encountered before' and this suggests they are being guided by rules.[160] Moreover, that they are applying the relevant rules to new cases can be ascertained explicitly, through acts of explanation and justification. Logical truths will either express these very rules or be derivations of them. This is not the case for firmly held, empirical beliefs, which by their nature cannot be derived from linguistic, logical, or mathematical rules. Indeed, Quine's scepticism regarding inexplicit rules or conventions is not found to be convincing by those who find his arguments otherwise compelling.[161] There are more general reasons for this. One can imagine circumstances where a convention has been adopted by a community even though no one has discussed it. There might be a mutual understanding, or expectation, for example, that students living together wash up after themselves. These rules may remain inexplicit for the duration of a tenancy, or they may be rendered explicit when someone fails to accord with them.[162]

So much for this objection, one might think. Yet some find it curious that Quine thinks the move to implicit convention can avoid his initial argument. Here is Bob Hale:

---

[157] Quine (1936 [1966], 97).
[158] Creath (2003, 248).
[159] Cf. Schroeder (2021, 98-100).
[160] Warren (2017, 131).
[161] See Hale (2013, 125).
[162] Lewis (1969) analyses conventions in terms of mutual expectations and argues they can be in force even when their followers are unable to articulate them.

> Resort to some inexplicit analogue of explicit stipulation – some sort of tacit communal agreement to treat the relevant sentences as immune to revision – would still run foul of the problem of finitude. For we would still have, at best, only finitely many direct guarantees of conventional truth, and any attempt to provide for the infinite remainder by appeal to logical consequence would serve only to highlight further as yet unconventionalized necessities (Dummett's point), or would degenerate into infinite regress (Quine's). (Hale 2013, 127)

This may be a problem for some accounts, though more plausible versions are surely immune from it. It may be that a basic rule formulation is: 'for all statements, if x, then y'. But the mistake Hale and others make is to think that basic rule-following consists in stipulating the general statement as true and then drawing inferences from it. Of course, if that were the case, they are quite right that an additional inference would be required from this universally quantified claim to a specific case. But that type of statement is merely a *formulation* of the rule.

We don't start by stipulating certain statements as unassailably true and then infer what follows from them. The point of moving to implicit inference rules is not merely to invoke an 'implicit surrogate or analogue'[163] of explicit stipulation.[164] The unassailable truth is a consequence of the fundamental role the norm thereby expressed has in shaping normative behaviour. It is that the technique we employ holds the relevant statements as true in all cases which accounts for their unassailable truth. To be in possession of inference rules is to have acquired an understanding and ability to react to various cases in the correct manner.[165] To learn rules is not merely a matter of learning to assign 'true' to certain sentences but is rather to acquire a technique of drawing inferences that can be expressed in rule-formulations like those given above. In logic, those rules determine the use of logical connectives and what counts as a proposition. The technique doesn't require us to move from the general rule to particular inferences, it is just that ability to draw those particular inferences (to use those connectives correctly). We might say the technique is expressed in, or stated by, the general formulation.

The techniques learnt can be used to derive further logical truths from accepted rules. In so doing, one does not appeal to any rules not already accepted. The drawing of logical consequences from rules is a matter of determining which other statements we must accept for

---

[163] Hale (2013, 126).
[164] A similar analysis to this is given in chapter eight of Nyseth (2018).
[165] Some might prefer to talk in terms of dispositions, but while our dispositions may give good evidence of the rules we are following, a focus on them rather obscures the fact that I can understand how to apply a rule without being remotely disposed to do so.

those rules to hold. The relevant proofs show what is required by the rules we've already adopted. In a similar vein, previously uncalculated sums have a determinate answer according to our mathematical operations.

An important connection can be made here with Wittgenstein's discussion of rule-following. Wittgenstein explores the possibility of rules being interpreted in many ways and the ensuing anxiety that rules are, therefore, powerless to determine their correct application.[166] Of course, this worry is exaggerated: just because something can be misinterpreted doesn't mean it has to be. But that is not sufficient to cure the anxiety, for we then face a potentially infinite regress. The correct interpretation of the rule does itself require a particular understanding for it to yield the right result. In other words, the correct interpretation of the rule is only correct when interpreted correctly, and likewise for the interpretation of that interpretation.[167] That this chain of reasoning is fallacious should be obvious from our ability to apply rules. For example, for 'x = 2y', x = 18 when y = 9. But what accounts for our halting the regress?

To help us, we should draw a distinction between rules and their formulations. Our understanding of the rule clearly cannot rely on our identifying the correct interpretation of the relevant formulation, for that path leads only to regress. Instead, we should recognise our ability to apply the formula as basic, unmediated by interpretations and other rules. We learn what counts as following the rule. A rule-formulation is not to be understood as an attempt to capture an ideal formulation, though some will no doubt do better than others. Rather, a satisfactory rule-formulation is merely a statement of the relevant rule. And what we need are not endless interpretations but an understanding of what the rule requires. We have little trouble seeing this once we have mastered a rule. But we are misled by the fact that rules can be misinterpreted. We wonder how it is that one can land on the correct interpretation of a rule without already understanding it. The truth, of course, is that we are not alone in such circumstances. We are trained to understand what rules require, not left to our own devices. We are shown what is correct and incorrect according to certain rules. These are the lessons from Wittgenstein's discussion of the rule-following paradox.[168] No formulation can stave off the possibility that it might be misinterpreted, but that is only a problem if we mistakenly believe that rule-following consists in being in possession of the perfect formulation. We know, however, that it does not. For the rules we follow are just those we use in explanations and justifications. At the most basic level, we learn to do something (to follow the rule). Of course, there are examples where we

[166] See Schroeder (2006, 185-201) for a detailed discussion of these passages in Wittgenstein.
[167] Cf. PI, §201.
[168] PI, §201.

learn a formula (or rule) via another, as in the case of "'y^2" is equivalent to "$y \times y$'". But, following Schroeder, we can say that:

> … such a translation of one formula into another cannot be the paradigmatic case of understanding, certainly not the basic case of understanding: or else we would never *begin* to learn the meaning of any formula. Likewise, our basic linguistic understanding cannot be accounted for in terms of translation into another language. (Schroeder 2006, 193)

We saw Hale suggest above that the move from general inference rules to specific cases is problematic because it requires us to make additional inferences from those rules. We have just seen an argument to suggest the move from rules to individual applications is problematic because we are required to make an infinity of interpretations. The answer to both is to move away from the picture of a perfect set of instructions, and instead recognise that our 'basic skills must stand on their own'.[169] These skills concern how to apply different rules. They can be developed through training and that training may include references to explicit rule-formulations. But those formulations cannot be the *basis* of our understanding. None of this is intended to exclude the many cases where we can and do explicitly introduce rules. Nor is it to suggest that on each of these occasions we start from scratch. The point, rather, is that what explains our ability to understand and apply rules without entering an infinite regress of interpretations are those basic, *practical* skills. Skills that are presupposed in our understanding of rule-formulations.

9. Conclusion

None of the objections considered, it seems to me, are compelling.[170] Now the account has been explained, argued for, and defended, we next consider whether language can support it. However neat an explanation of necessity the normative account provides, it relies on linguistic explanations. A convincing defence of the account must, therefore, show that language is a suitable source of necessity. Is language such that modal truths can be explained by way of it? It is to this question that I now turn.

---

[169] Schroeder (2006, 194). Cf. Wright (2008).

[170] There are, of course, many other objections that have been discussed and challenged elsewhere by those sympathetic to the normative account. For example, I have said little about logic and even less about mathematics. For logic, see Nyseth (2018) where several objections are explored. On mathematics, see Floyd (2000) and Schroeder (2011; 2014; 2015). Baker and Hacker (2014) and Warren (2015b; 2020) consider objections from both logic and mathematics (amongst other things). And Schroeder (2009) considers some objections relating to necessities expressed more readily in natural language.

Part II

The normative account of necessity presupposes a distinction between conceptual and empirical truths. The view is that 'necessary truths' are not accurate descriptions of the world but instead reflect conceptual norms that make descriptions possible. This is controversial.

Several philosophers, particularly those influenced by Quine, have argued the distinction is untenable. They argue we cannot identify meaning-determining principles as separate from descriptions of the world. Language manifests no such distinction. If this were the case, necessity could not be explained by appeal to conceptual propositions. Moreover, if there were really no distinction between conceptual and empirical propositions, then it could not be said that conceptual truths do not answer to reality. For they could not be adequately distinguished from empirical ones, which indisputably *do* answer to reality.

One way this problem is sometimes put is to say language is not as simple as the distinction between empirical and conceptual propositions supposes. There are no simple rules to read off and so determining whether a given proposition expresses, or follows from, a linguistic norm is not at all straightforward. This might seem particularly acute for my account because, as we will see, I think necessity is an essential aspect of language with descriptive power. So, one might further object that, even if there were linguistic norms, they would be so complicated and admit of enough exceptions that they couldn't possibly be characterised as 'necessary truths' when stated.

The answer to these challenges, which amount to the claim that language's character is such that the normative account of necessity cannot be sustained, is to locate necessity in the relation between norms and concepts. It is not a condition on the content of norms. This means that the content of norms can be sufficiently complex that in some cases there may be no determinate answer as to whether a proposition is conceptual or empirical. But this does not undermine the in-principle distinction between the conceptual and empirical. It merely reflects what the norms are, and the character of the concepts involved.

(5) Necessity, Exceptionlessness, and Generality

In what follows, I describe different ways in which norms of representation might defy a simplistic picture of language constituted by unqualified necessities. The central claim is that a more complex account of language can be maintained while making room for constitutive relations (the likes of which are articulated by 'necessary truths'). This makes it possible to see how an in-principle distinction between conceptual and empirical propositions can survive genuine indeterminacy in particular cases. Insofar as it can be unclear whether a given feature is necessary, for example, a proposition stating that the object has that feature is unable to be characterised as either a conceptual or empirical claim. This point is developed in the next chapter, which considers in detail some arguments that suggest the more complex account of language is incompatible with the kinds of meaning-determining principles and constitutive relations my account relies upon.

1. Exceptions

I have argued that necessary truths are expressions of conceptual norms, often in the form of rules for the use of words, that are unconditionally true. The firmness of necessary truths, such as 'all grandfathers are fathers', is the result of the relation that holds between the norm and the concepts contained within it. It is that something *functions* as a norm that accounts for the firmness. That a grandfather must be a father is a consequence of those concepts being defined as they are, of their being used in accordance with that norm. If they were used differently, and that firmness was softened, the words 'father' and 'grandfather' would have their meanings changed. How great a change that would be depends on the extent to which the norms were amended.

But while the function of a proposition, its role in language, may make it clear that it is being used as a norm of representation, this places little restriction on the content of that norm. Thus, while norms like the above, typically cast as 'necessary truths' in philosophical circles, qualify, there is no reason why norms could not appear altogether less firm. Statements such as 'a grandfather must be a father' grant no exceptions. The norm appears general and – in a sense – straightforward, without qualification. There is, however, no reason to assume that statements that perform similar functions must themselves manifest the same characteristics. A norm's function provides the power of the 'must', but it does not necessitate the kind of exceptionless nature that typifies 'necessary truths'. If the norm had been 'a grandfather must be a father

except when his name is '"Ted"', we would be less inclined to describe it as a 'necessary truth'. For what we call 'necessary truths' typically exemplify absolute generality. Likewise, we would not be inclined to describe the above, qualified statement as comprising the essence of grandfatherhood (more likely to say it lacks an essence). However, we could still say: necessarily, a grandfather is a father or named 'Ted'.

The point, then, is that necessity resides in the constitutive relation that holds between a norm and concept, but this places no constraint on the *content* of that norm. While an exceptionlessness relation holds between the norm and concept, the norm itself can grant exceptions. One way of expressing this would be: there are no exceptions to the norm beyond those already granted by it. Once those exceptions are accounted for, the statement holds in all circumstances. We might, on reflection, conclude that the notion of there being an exception here is itself confused (or at least confusing). For the 'exception' is part of the norm, not an exception to it. This is well put by Frederick Schauer in his paper 'Exceptions', which deals with similar problems in the law. He writes, regarding the rules of baseball:

> . . . we would ordinarily expect exceptions to be built into the meaning of a primary technical term. Because foul balls are not home runs in the first place, it is odd to say that foul balls are exceptions to the rule defining home runs. (Schauer, 1991, 877)

It is indeed odd in that case, and perhaps many others. But where the easiest explanation of a term requires one to state what would be a simpler rule combined with one or two exclusions, it seems more natural to describe them as 'exceptions'. Perhaps this contrast has something to do with the way baseball is defined according to strict rules, and taught as such, whereas when studying language we may be describing something that has developed more naturally, or at least that we were not taught according to strict rules. Exceptions are ubiquitous in our descriptions of language, not least because we so often describe it *ex post*.[171] Our describing it *ex post* means we seek regularities in our correct usage that may nevertheless fail to cover all uses. We find principles that hold in most cases and then exclude the few cases where they do not. We are already engaged in the practice of language before we go about deciphering some of the norms by which we regulate our speech, since we do not tend to learn language via strict rules or definitions. Moreover, it is rare for us to consult rules when using it. As Schroeder writes:

[171] Cf. PI, §31.

> … the piecemeal normativity of language may sometimes be conveniently described in terms of rules, but speakers do not actually follow these rules. They figure in *ex post* descriptions of language, not in its actual practice. (Schroeder 2013, 163)

Thus, the presence of exceptions in our descriptions of language, our inclination to describe them as such, is in part a reflection of our describing language *ex post*. It reflects our attempt to provide guides to correct word-use, while acknowledging that simple, unqualified rules often fail to do justice to its complexity. There are other reasons, of course. One might be that language develops naturally and may involve granting certain exceptions over time that are like cases we already accept or reject. Likewise, our propensity to speak loosely, figuratively, and more generally project words into new circumstances may eventually erode what were once stricter, or more confined, uses of terms. This explains why Wittgenstein asks us to…

> … remember that in general we don't use language according to strict rules – it hasn't been taught [to] us by means of strict rules, either. We, in our discussions on the other hand, constantly compare language with a calculus proceeding according to exact rules. (BB, 25; cf. PI §81)

One example of this phenomenon concerns a favourite of philosophers. Marcus Giaquinto suggests a plausible counterexample to 'a bachelor is an unmarried man'.

> Peter once married a refugee merely to save her from deportation to a country suffering civil war; after the marriage ceremony he never saw her again, but continued to live as young single men do, available for a long-term personal partnership. In the eyes of those whose idiolects are now under discussion, Peter remained a bachelor; beyond legal contexts it would be wrong to deny it. (At the same time it would be accepted, without any inconsistency, that there is a legal use of "bachelor" under which Peter was no longer a bachelor, just as there is a legal use of "guilty" according to which a murderer is not guilty of the murder if he has not been declared so as a result of a legal trial.) (Giaquinto 2008, 97)

This passage suggests there is a legal use of the word 'bachelor' that applies standards that are otherwise redundant in other uses of the term. Giaquinto makes a plausible case for there being a non-legal use that is not so restrictive, and that would permit Peter being counted as a 'bachelor'. His second example is the inverse of the first: Mike is in a committed relationship with his partner with whom he has two children. They have lived together for several decades. We could go further than Giaquinto does and suppose that Mike even believes he is married

when he is not, perhaps because of a mistake in the registry office. The suggestion is that Peter is a married bachelor, while Mike is an unmarried non-bachelor.

Giaquinto's cases are plausible. The use of 'bachelor', as in 'bachelor's pad', for example, is not obviously restricted only to unmarried men. It often seems directed at people exhibiting certain forms of behaviour (no doubt behaviours often associated with unmarried men). Supposing that Giaquinto is right about this non-legal use of the word 'bachelor', we hardly have a formula for determining what the relevant exceptions are. On the contrary, it seems like a rough estimation of the word's use would be something like 'bachelors are unmarried men except in circumstances where a man's marital relationship is especially thin (for example, in cases where it is really a legal relationship and no more), or where an unmarried man remains in a committed relationship for some time.' These exceptions are not terribly precise, and they may turn out to be too permissive or restrictive. Even if this guide to this use of 'bachelor' is correct, there will likely be borderline cases that lack a definitive verdict. But such vague exceptions do not threaten the linguistic account of necessity. They merely reflect the content of the norm and so the shape of the concept. The constitutive relation remains. It just turns out there is a concept 'bachelor' that is not as precisely defined as alternative uses of the term. Notice that there will still be 'necessary truths' that follow from this norm. For instance, a man cannot be a bachelor if he is married to, and in a committed relationship with, his wife. The non-legal use of the term 'bachelor' was bound to become less clear-cut as social norms regarding, for example, cohabitation changed. It may over time fall out of use altogether. Norms often cover only the cases we expect. Social or technological changes may produce situations where it is no longer clear how they apply.[172] Just as emerging technologies, for example the internet, raise novel questions for copyright law which is adapted for the changing circumstances.

There are other cases where we would be disinclined to talk in terms of 'exceptions', where the use of some word is so varied that one must provide the context in which it is used for a reasonable answer to be given as to what it means.[173] Waismann suggests 'time'.[174] One can sometimes substitute 'time' for 'opportunity' in 'now is your time', but it would be absurd to do so when I say, 'the time is five o'clock'. And 'in no time at all' and 'time after time' it continues to resist general treatment. Contrast this kind of case to another one of Waismann's examples where we might be inclined to talk of there being an exception to the rule. Normally, one cannot

[172] This is discussed further at the end of this chapter.
[173] Travis (2006) contains some examples of this kind.
[174] Waismann (1950, 26-7).

see something without looking at it, but in the case of after-images, this does not hold.[175] I can neither look at an after-image nor look away from it. Seeing after-images is thus different from seeing external objects. After-images provide an exception to what is otherwise a relatively general rule.[176]

Returning to the 'must', it does indeed introduce a kind of exceptionlessness, insofar as whatever qualifications exist within the rule, they are nonetheless part of that rule. And that rule must be followed for one to successfully use the relevant word with its established meaning (for to use it in accordance with that meaning simply is to use it according to that rule). We can understand this as a restatement of the observation that these rules are constitutive and therefore a concept does not take its particular shape unless used according to the patterns of usage that give it that shape.

This – in a sense – brings us into agreement with the primary point Schauer made earlier, namely that what I am calling 'exceptions' are in fact part of the rule. While 'exception' was helpful to get the idea going, perhaps 'qualification' and 'unqualified generality' are better. Now that the matter has been clarified, my future usage should be clear enough. I am not seeking to dispute the point Schauer makes, though I may, on occasion, drift from him terminologically. My willingness to do so stems from the fact that I am taking for granted that the rules I am interested in are constitutive. Any exceptions I speak of are, therefore, part of the rules under investigation, and thus help to characterise the relevant concepts. For now, I am assuming that the exceptions made to a rule can be formulated explicitly, but I do not assume that is true of all cases. Perhaps where such exceptions cannot be captured by a rule, we might want to call them exceptions *to* the rule. That description is not without its problems, however, for those exceptions still help form the grammar those rules (partially) capture. The rule could hardly be characterised as constituting the concept if there was not some way of acknowledging that it provides only a simplification of correct word-use. As Waismann writes:

> The philosopher wants to recover the regular use for [a word]; he wants to fill up the great gaps left by normal usage: and yet these gaps are just what is characteristic of the meaning of such words. (Waismann 1965, 84)

[175] Waismann (1951a, 58).
[176] Other exceptions might include certain metaphorical, or figurative, uses of different terms. Whether there are rules that govern specific metaphorical uses is a moot point. It may just be that we accept a general norm that allows for metaphorical uses of many words. Cf. Mulhall (2009b, 158-161).

The rule left unqualified is not constitutive of the relevant concept at all, for it draws boundaries that include the exceptions that we reject as incorrect uses (or excludes exceptions that we accept as correct uses). Nevertheless, an unqualified rule may prove a helpful guide to correct word-use.

2. Checklist guides

So far, the lesson has been: a proposition's function as a norm is what makes it constitutive of the shape of the concepts contained within it, but the shape of those concepts is determined by the content of the rules that pertain to them. This can be seen again in cases where we do not have rules so much as lists of criteria to be fulfilled.

For instance, I might lay out several criteria necessary for someone to count as a 'good friend'. These might include qualities that a 'good friend' has and features of our relationship. But I do not need a rule of the kind: 'every one of these criteria must be met to a certain degree for them to count as a good friend'. Instead, it might be flexible: there might be some qualities that are essential (i.e., criteria that must be met, criteria without which one would not count as a 'good friend'), and others that, while some undetermined number must be met, there is no rule specifying which. Among the latter, different weights may apply such that a smaller number of the weightier (or more significant) criteria need to be in place for someone to so count. It seems unnatural to describe anything other than those criteria that must be met as essential, and yet some undetermined number of the other criteria *must* be before one could be counted as a good friend. Both the essential criteria and the others constitute what I would be willing to count as a 'good friend'. Any exposition of that concept requires a description of all the criteria, both those that are individually necessary (or essential) and those that are not. Importantly, we have no reason to think there are individually necessary criteria for every concept. We might have a set of criteria where none are individually necessary, but several of them must be present for some object to qualify as being an instantiation of that concept. To sum up, we might say: that something is a norm means that it determines what a 'good friend' is, but what counts as a 'good friend' (i.e., what the qualities are and how they need to be combined) is down to the norm's content.

We can imagine more explicit uses of checklists, in the evaluation of a company's employees, for example. There might exist a set of criteria that determine employees' level of performance. Someone counts as a 'high performer' if they meet a certain number of the relevant criteria, which might detail (among other things) their contribution to the company's

revenue through services they provide and their generation of new streams of revenue. This in turn might become important when deciding who should be promoted. Checklists of this kind are quite common across language and can typically be identified by the types of explanations we provide. Often when asked what we mean by something, we will give an inexhaustive list of criteria. Examples in which such explanations might be provided include 'stable government', 'functioning democracy' and 'healthy ecosystem'.

What I appear to be saying, then, is that (in some cases) there are still necessary or essential criteria (or conditions), but such criteria do not exhaust the conceptual realm. For the conceptual realm consists of all that is covered by the rule, including those things that, while (individually) inessential, play an indispensable role in determining the shape and character of a given concept. Thus, while essence might be absorbed into grammar,[177] essence does not provide a complete characterisation of all that is grammatical. But insofar as rules are constitutive of the concepts contained within them, we should add that they will be expressible by a 'must'. For once we have the rule, qualified or not, ready to be formulated, the 'must' will signify that something simply would not instantiate that concept without satisfying that rule.[178] And that is merely a condition of what it takes to be constitutive.[179] Conversely, where something *does* satisfy the rules, it follows that it instantiates the relevant concept.

One might worry here that I am betraying an argument I made against Kalhat. There I accounted for necessity by appeal to the rule's function. Yet now I appear to be reintroducing the notion of necessity into my discussion of criteria while maintaining a distinction between the content of a rule (what those criteria are) and the rule's function (what makes them criteria). Is the necessity of those individual criteria left unaccounted for? This worry can be assuaged by the observation that necessary criteria can be formulated in rules such as 'An X must have A to be an X'. The necessity of those criteria is still explained by the constitutive relation between the rule and the concepts contained within it. The 'must' remains a result of conceptual norms (constitutive relations). We can talk of 'necessary criteria' precisely because those criteria bear that constitutive relation, namely 'X would not be an X without A'. With respect to (individually) unnecessary criteria, this relation does not hold. The relevant relation exists between a range of

[177] 'Essence is expressed by grammar' (PI, §331).
[178] Cf. 'The mathematical Must is only another expression of the fact that mathematics forms concepts.' (RFM 430)
[179] I think this is what Forster (2004, 12) is getting at when he says: 'it is easy enough to see that these further parts of grammar can plausibly be considered to possess (non-causal) necessity'. By 'further parts of grammar', he means statements that do not share the character of so-called 'necessary truths'. This point also holds of family-resemblance concepts. There we cannot say: 'these features must hold if something is to count as. . .', but we can say: 'for something to count as. . . it must be a member of this family'. Again, this is merely what it takes for something to be constitutive. Something instantiates a family-resemblance concept in virtue of its resemblance to other members of that family.

disjunctive criteria and the thing in question, and thus it cannot be said that individually those criteria are necessary.

3. Family-resemblance

At §67 of the *Investigations*, Wittgenstein introduces his now famous family-resemblance concepts.[180] These concepts are distinct in that they appear to lack discernible rules for their application. This is not merely a matter of our struggling to describe the regularity by which we use these words via clear rules. It is also a response to the kinds of corrections and explanations we give. The distinction between family-resemblance concepts and checklist guides is unclear and certainly not sharp. Perhaps family-resemblance is distinctive in the sense that explanations tend towards examples rather than criteria. Typically, family resemblance concepts will be explained via a set of examples that resemble one another, and that have come to be grouped naturally over time. Rarely does there seem to be an explicit corrective to misuses, other than a mere assertion that something does not belong to that category. One of the defining features of family-resemblance concepts is that a given member of a family need not share a resemblance with every other member of that family. Instead, it might only be connected to that member via a third, where each resemble the third member in ways they do not resemble each other. A helpful parallel to such concepts can be found in genetics. That you inherit genes from your parents is a given, and that your parents inherit some from your grandparents is likewise a given. But because a child inherits some genes from one parent and some from the other, it is possible (and likelier the further one goes back) that after many generations, a descendant will, in fact, find ancestors from whom they bear no DNA. Nevertheless, there are threads throughout a family tree that will link them back to that ancestor via inherited genes of common family members.

Wittgenstein's primary example of a family-resemblance concept is 'game'.[181] He suggests there is no common feature across all games, but rather they resemble each other in important ways. This example has been criticised.[182] Whether the criticism is fair I shall not attempt to answer here, but it is worth pointing out that it often risks missing the point.[183] That one can, in fact, develop a rule that fits the pattern of our use of some concept does not show that the

[180] A good discussions of this can be found in Mulhall (2001, 81-87). Mulhall counters criticisms made by Rundle (1990).
[181] Waldron (1994, 518) suggests 'religion'. Another might be 'sport'.
[182] Suits (1978). Schroeder (2006) presents convincing counterarguments.
[183] Schroeder (2006) makes this argument.

concept is constituted by that rule. For while the rule may work for current applications, it may later be shown to be an inadequate description of word-use by our including an example that would be excluded were we following the proposed rule. Coming to understand the rules of representation that are in use is not merely a matter of looking at patterns in our applications.[184] How we correct people, what we say when we correct them, and how we explain the meanings of our terms are all important features in determining what norms we are following. When critics attack Wittgenstein's example of 'game' by claiming they have discovered a formula that fits our application, we may still conclude that the rule does not govern our application. For it is often far from obvious any of us were, or could have been, following that rule, given the kinds of explanations we provide. Moreover, that the discovery of such a rule can come as a surprise provides some evidence against it being a rule of ours. Linguistic reminders ought to be exactly that. Whatever we think of potential examples of family-resemblance concepts, all I argue here is that a given concept *could* take this shape. That it is possible for a concept to develop not dissimilarly to the way genetic material is inherited through a family tree. Our not appealing to rules in the explanation of what a game is supports this possibility.

Finally, it is worth bearing in mind that family-resemblance may be an aspect of a concept that is also characterised by some of the other features explored in this chapter. It seems perfectly possible, for example, that we might have what is a very clear rule, but for a family of exceptions. For many, I would suggest their experience of language regarding the natural world shares important similarities with this already. When I am told that an orca is, in fact, a dolphin, or a mountain goat is not really a goat, it is unnecessary for me to learn why that is the case. I can take the zoologist's word and accept that my understanding of the relevant concepts is not as advanced as theirs, thereby accepting that my concept is theirs. It is a feature of our culture that this kind of deference exists, that we do not have (in all areas) lay-concepts and expert-concepts, but rather accept that some of us have lesser understandings and therefore stand to be corrected.[185] For the most part, it makes little difference to our lives. If we can accept exceptions based on presumed authority without our word-use becoming hopelessly disordered, I take it that we could easily have rules that admit exceptions that share some resemblance we grow sensitive to throughout our training. Indeed, it seems our language could survive seemingly random exceptions, which for some is how the zoologist's interventions are effectively received.

[184] Waismann (1965, 129-152) explores how we ascertain the linguistic norms we follow, and the inherent difficulties of doing so, in a systematic fashion.
[185] Lassiter (2008, 621-622) acknowledges this point and credits Putnam (1975) for it.

A degree of regularity is required for distinguishing correct applications from incorrect ones, but beyond that there is a vast sea of possibilities.

Here again, necessity is not absent or compromised by virtue of the way these words are employed. Something instantiates a given family-resemblance concept by virtue of its belonging to the relevant family. Given it belongs to that family, it must instantiate that concept. This is once again a reflection of the fact that the concept is constituted by the relevant norms of use (which in this case are a matter of family-resemblances).

4. Vagueness[186]

Vague concepts are those that have ill-defined or imprecise boundaries.[187] Some applications are therefore subject to genuine indeterminacy. There exist borderline cases 'in which one just does not know whether to apply the expression or withhold it, and one's not knowing is not due to ignorance of the facts'.[188] The vagueness of a concept does not render all applications indeterminate.[189] It may be very clear that the sky is blue, but less clear whether someone's turquoise shirt is. It is obvious that Gordon Brown was a politician, but less so that every local councillor is. That a concept is vague does not mean that it has no clearly defined boundaries. 'Blue' may be vague when it comes to distinguishing it from certain shades of green, but that blue is a colour is not uncertain.

Vagueness is distinct from ambiguity. The latter arises when the same word is used to express different things.[190] When I talk of approaching the bank, it is ambiguous as to whether I mean the financial institution or the edge of a river. Vagueness is also distinct from specificity. The concept 'tree' is more specific than the concept 'living thing', 'but it may be much easier to say whether a given object is a living thing (as opposed to inanimate) than whether it is a tree (as opposed, for example, to a bush).'[191]

---

[186] Much has been written on vagueness. Keefe and Smith (1996) includes an informative introduction and several influential articles on the topic. Sorensen (2018) and Hyde and Raffman (2018) provide surveys of the literature. Influential books on the topic include Williamson (1994), Keefe (2000) and Shapiro (2006).
[187] Poschner (2011, 139-140) gives five reasons for our language exhibiting this kind of vagueness.
[188] Grice (1989, 177).
[189] I take this to be obvious, but it has been denied. Ludwig and Ray (2002, 421) claim that 'semantically vague predicates neither apply nor fail-to-apply to anything'.
[190] Sorensen (2001, 413) provides a more technical distinction: ambiguity concerns words and sentences (their pre-propositional forms), while vagueness concerns the concepts and propositions those words and sentences are used to express. In the language of the *Tractatus*, ambiguity concerns cases where we have one sign playing the role of multiple symbols. For the author of the *Tractatus*, the ideal language would remove such ambiguity by ensuring every symbol has a unique sign (3.323-3.325).
[191] Waldron (1994, 522).

Vagueness is a common feature of our concepts, and it comes in degrees. The logical rules of inference, for instance, do not seem vague. The rules of chess, as they are now, seem to be very clear as to what defines a pawn and the moves it can make in a game of chess. Likewise, elementary rules of mathematics do not seem to provide much scope for vagueness. To be sure, it might be unclear whether I should count one or two people when faced with certain conjoined twins, but that is a result of the indeterminacy around what we are here counting as a person. We would need more details of the case to have any prospect of answering, and we ought to be prepared for the possibility that our current understanding of personal identity admits of no determinate answer in some cases (suppose, for example, they had one brain but, at times, two distinct personalities). Moreover, precisely what would follow from their being so counted may not be obvious. Many ethical questions, for example, are bound up with our concept of personal identity, so unusual cases can prove disruptive. However, if the reason for my counting is that I am trying to determine how many chairs I need to arrange, and the conjoined twins share all their body but for their necks and heads, then I should count them as one (though I may not know how many meals to serve). Context has the potential to eradicate vagueness, but it is no guarantee.

Jeremy Waldron suggests that vagueness can 'arise whenever we confront a continuum with terminology that has, or aspires to have, a bivalent logic',[192] such as in the case of the colour spectrum. But it does not need the presence of a continuum to arise. It is far from clear that there is a spectrum of tablehood, yet it might be unclear whether we should say a chest that is used as a coffee table really counts as one. We could understand people who gave conflicting answers on such matters, and few of us would maintain that there exists a definitive one. Vagueness here is sometimes a result of what Ralf Poscher calls our 'indecision as to certain necessary properties'.[193] But as I have argued above, there is no reason why it should be kept to 'necessary properties', unless this is also intended to cover those criteria that may be individually inessential. William Alston's combinatory vagueness is a version of what I have in mind, whereby it is undecided how many relevant characteristics something is required to have in order to qualify as the object in question.[194] We might add that, where those characteristics come in degrees, it might be undecided the extent to which different characteristics need to be present

[192] Waldron (1994, 516). This covers the vagueness of individuation too, i.e. mereological questions about where one object stops and another begins (a result of reality's being continuous). This kind of vagueness is pervasive, see Hyde (2008, 112).
[193] Poscher (2011, 132).
[194] Alston (1967, 219-220).

for something to so count. And it might be undecided precisely which combinations of these different variables can make an object so count.

One implication of this is that vague utterances may say nothing (true or false) at all.[195] Consider 'child'. The idea is that some cases are genuinely indeterminate, meaning we can neither say that 'child' applies nor fails to apply to a borderline case. This does not mean we could not understand someone who claimed it to be true of a borderline case. The context might make it clear that their application is correct. Where there is no pre-existing context, one can imagine saying something like: 'you can choose to count them as a child if you like, it doesn't really matter'. Moreover, we could happily understand two people, in the same conversation, choosing to count the same person differently. In such cases, it is acknowledged that there is a degree of arbitrariness in their decisions, which distinguishes them from cases where we are determining whether certain laws pertaining to 'children', meaning 'persons under eighteen', apply. But where this additional context is not given, the concept is not sharply defined enough to provide a definitive answer. Thus, unless the context makes it clear, we could say one of three things: that they are a child, that they are not a child, or that they are neither a child nor not a child (a refusal to be drawn on the issue, the acknowledgement of a borderline case).[196] The former options are available because we accept speakers using their discretion in many contexts where standards are not sharply defined.[197] The latter is fine insofar as we accept those standards are not sharply defined.

This is not without controversy. Indeed, Williamson argues that it is incoherent. It must be either true or false that someone is a child. If we say they are not a child, then it must be false. There is no logical space for the suggestion that they are neither a child nor not a child. But notice how this way of presenting things obscures the point just made. When refusing to be drawn on the issue, one does not say that the person in question is not a child (i.e., that it is false to say so), but rather that 'true' and 'false' do not apply here because the concept provides no determinate answer. We are not able to rule them in or out. Williamson offers three arguments in favour of his position.[198] It has been suggested that two of them beg the question.[199] The third, however, informs our discussion of vagueness and is not obviously question-begging. In effect, Williamson argues that whatever we have to say about borderline cases, we must accept the following statement (and many of its kind): 'if the borderline case is a child, then its younger

[195] See Moore (2019, 50-52).
[196] See Hanfling (2001, 35) for why this needn't commit us to a third truth-value.
[197] Cf. Shapiro (2006).
[198] Williamson (1994, 195–197).
[199] Moore (2019, 51-52).

sibling must be a child'.[200] Furthermore, for that conditional to be true, so he argues, it must say something. And for it to say something requires both its antecedent and consequent to say something. Moore's counter to this line of thinking is compelling:

> The person making this utterance is stating a rule of representation: a prohibition against counting the addressee [the borderline case] as a child without also counting any younger sibling of theirs as a child. But accepting this rule is quite compatible with giving each of various individual verdicts on the matter. In particular, I cannot see how it precludes refusing to count the addressee as a child and refusing to count the addressee as not a child. The truth of the conditional (that is, the holding of the rule) allows for, among other things, the lack of truth or falsity of an isolated appearance of its antecedent.[201] (Moore 2019, 52)

Sober offers another example. He suggests that even though 'rich' and 'poor' may be vague and heavily context-dependent, the rule 'someone who is poor has fewer resources than someone who is rich' holds come-what-may.[202] It may be difficult to determine whether someone counts as 'rich', but if I am to count someone as 'rich', they'd better have more resources than those I will count as 'poor'. That this internal relation holds between 'rich' and 'poor' means that what counts as 'poor' adjusts with what counts as 'rich' in a given context.

Williamson is, of course, a sceptic of this kind of vagueness.[203] He thinks vagueness is a result of our epistemic imperfection, that from a God's-eye perspective one could determine where the conceptual boundaries really lie. While I have not the space to answer him in full, the arguments I have made elsewhere should make clear the kind of response I would make. The strangeness of Williamson's account resides in the thought that there is something to determine such sharply defined boundaries. In a paper criticising Williamson's 'Master Argument' for epistemic vagueness, Greg Ray pulls out the following quotations.[204]

---

[200] Williamson (1994, 196).
[201] It is worth flagging here something Moore points out in a footnote of his own (nn. 54). One might think Williamson forestalls this objection by talking about 'material' conditionals (Williamson (1994, 196), where a material conditional's truth is determined by the truth and falsity of the antecedent and consequent. As Moore points out, however, it is question-begging to suppose there are any relevant conditionals that are both material and true (or material and false).
[202] Sober (2000, 245). Russell (2008, 179) also discusses this.
[203] Enoch (2007) relatedly argues that the distinction between nihilists and epistemecists is semantic rather than ontological. What they disagree on is the best understanding of vague discourse. Both agree there are no vague concepts in the standard sense.
[204] Ray (2004, 188).

> You have no way of making your use of a concept on a particular occasion perfectly sensitive to your overall pattern of use, for you have no way of surveying that pattern in all its details. (Williamson 1994, 232)
>
> A small shift in the distribution of uses [in the community] would not carry every individual use with it. (*Ibid*)
>
> Vagueness is a source of inexactness, . . . because individual uses of a vague term are not fully sensitive to small differences in the overall pattern on which small differences in meaning supervene. (*Ibid*, 234)

Ray's counterargument is that there is nothing, in principle, that prevents us from being sensitive to our overall patterns of use. I would go further and say that in many cases there is no evidence that our terms are determined in as precise a way as Williamson supposes. Take the example of baldness. That we do not count the number of hairs on someone's head to determine, for certain, whether someone is bald is good evidence for our concept not being sensitive to the precise number of hairs on someone's head. That competent speakers can disagree over particular applications, despite being in possession of all the same facts, is further evidence that the concept is vague. Moreover, explanations that we tend to give of that concept do not support the suggestion that its conceptual boundaries are themselves as well-defined as some sceptics of vagueness would have us believe. And even if it was revealed that we all, in fact, assent to saying someone is bald at the same number of hairs on someone's head when presented with a sequence of balding heads, it would not prove that the concept was prepared for such precision. How might we ascertain this? Well, through the responses people give. It is not uncommon for us to recognise borderline cases and say things like: 'well I would count them as "bald" but I'm not sure there is a definitive answer. I think you could say either'. That is, we expressly treat our concepts as vague. Moreover, we would not accept the tallying up of our individual applications as a way of determining *exactly* where the line of baldness is, because we do not use that term according to that kind of precision. How could one interpret the procedure as giving us that kind of result when those making the judgements were not counting hairs? To engage in such a procedure in the first place presupposes that there are standards of correct use in play. Whatever those standards are, clearly they do not concern the precise number of hairs on any given head. For if they did, the procedure would be rendered pointless: if correct application is sensitive to the number of individual hairs on one's head, competent users would need to count those hairs to be certain they are correct. Yet none of us would. Whatever concept 'baldness' Williamson is working with, it is not one with which we are familiar.

Concerns about vagueness typically arise from the Sorites Paradox.[205] The paradox's power resides in how it seems to do justice to the vagueness of words like 'heap' and 'bald', while trading on the presupposition that they are more precise than they really are. Those concepts are not fit for the fine-grained distinctions the presupposition requires.[206] An argument typical of the paradox is as follows.

i. One grain of sand does not make a heap.
ii. For any number n, if n grains of sand are not enough to make a heap, then n+1 grains of sand do not make a heap.
iii. Therefore, there is no number of grains that can make a heap.

The second premise appears to account for the vagueness of 'heap'. Yet we can reject this premise and still refuse to affirm its opposite: namely, that there is a number n such that n+1 constitutes the least number of grains necessary to make a heap. A presupposition of both is that our concept 'heap' is prepared for such granular distinctions – as though counting grains of sand were a way of determining whether something is a heap as opposed to it being a matter for casual observation.[207] But I find no compelling reason to accept this, which is to say no reason for thinking either statement acts as a rule of representation. Whereas I have every reason, for example, to think 'a single grain of sand cannot make a heap' does.[208]

5. Abnormal circumstances, unanticipated cases, and open texture

To some extent, we met 'open texture' in the previous section. Some of our concepts do not appear to be vague until a set of unusual circumstances arise. These unusual circumstances exploit boundaries of our concepts that are otherwise left unchallenged by the normal conditions of their application. But an unanticipated case rendering an otherwise clear concept vague does not exhaust what Waismann calls the 'open texture' of a concept, which he says is 'something like *possibility of vagueness*'.[209] A concept's open texture is an acknowledgement of the possibility that our existing understanding of a concept may be challenged by an unusual case. This

[205] A discussion of Wittgenstein's approach to vagueness can be found in Faulkner (2010). She focuses on *Philosophical Grammar*, and the model Wittgenstein proposes as a way of thinking about vague concepts.
[206] Russell (1923) makes a similar point, describing the 'penumbra' in which our concepts are inapplicable due to their not being prepared for such distinctions.
[207] I think Faulkner (2010, 177) is suggesting something like this when she says the paradox 'treats the concept of a heap as if it functioned in a way that rests on its component grains.' The trouble with formulating it like this is that something's being a heap does, of course, rest on its component grains, just not in the granular way sceptics of vagueness think.
[208] For a handling of the paradox I find convincing, see Moore (2012, 590-593). See also Hanfling (2001).
[209] Waismann (1945, 123). See Shapiro and Roberts (2019) for a brief discussion.

challenge can manifest as our existing criteria clearly failing to settle the matter (thereby exploiting the vagueness of a concept's boundary that under ordinary circumstances is sharp), but it does not have to. Sometimes our criteria may seem to suggest the 'wrong' answer, pressing us to reconsider our current criteria. One explanation for this could be that a given concept has a particular purpose that is satisfied by its existing criteria in ordinary circumstances, but later comes to be undermined by unanticipated cases.[210]

Let us consider the first kind, where our concept suggests no answer for the unanticipated case. In these circumstances, we have several options. We can decide how to characterise the phenomenon with our existing concept, thereby stipulating a new boundary for it; we can develop a separate concept to account for the phenomenon; or we can accept the phenomenon as a borderline case and say it is neither correct nor incorrect to apply our concept to it. Consider this example. It is generally clear whether someone is younger than eighteen years old, but this clarity depends on one's date of birth being clear. Humans are only born once. Now, suppose a single person had to be reborn four times across several weeks. That is, they fully emerge, perhaps for some time, only to re-enter their mother and for the birth process to begin once more. How do we determine what their date of birth is? For legal reasons, it seems as though we would at least need to give one definitive answer. But that is only one context among many and may be subject to challenge periodically. It is at least possible that in other contexts we would not wish to stick to the answer that applies in legal contexts. Wittgenstein's disappearing chair is another example.[211] Given the criteria we have for chairs, it is not at all clear whether a chair that has a habit of vanishing does itself count as a 'chair', in part, I take it, because it is unclear whether we should count it as a physical object. Clearly it seems to be when it has not vanished, but that it vanishes seems to suggest it fails to persist. Is it an illusion? And yet it returns, perhaps in the same spot and in the same condition.[212]

A second type of case is where we have criteria that seem to settle the matter, but that do not appear to give (unambiguously) the correct answer.[213] Suppose someone dies but days later emerges as a zombie. Were we to examine their internal organs, we would find them be

---

[210] Waldron (1994, 537) identifies these kinds of unanticipated cases in the law, where a law clearly rules something out that it was never anticipated to ('no vehicles in the park' was never meant to prevent the emergency services from entering, for example). Open texture was made prominent in legal philosophy by Hart (1961). For a discussion of Hart's conception of open texture as it relates to law, see Bix (1991). For a critical discussion, see Schauer (2011).
[211] PI, §80.
[212] Cf. Z §350
[213] Waismann's (1945, 122-3; 1965, 223) form of gold might be one example. If it has the same atomic number, then it might seem to qualify as gold on existing criteria, yet we may feel it necessary to amend our concepts so as to take into account that it emits a novel type of radiation. Austin (1961) captures this phenomenon with his example of something that meets the definition of a 'goldfinch' but later explodes or quotes Virginia Woolf.

consistent with that of a dead person. And yet zombies do not resemble the dead in important ways: they continue to move around (seemingly of their own volition) and feed (albeit with different eating habits).[214] Often when zombies are discussed in popular culture, they are categorised as the 'living dead', a way of capturing their set of paradoxical features. The 'living dead' is a category that is quite apart from our usual ones, and it speaks to a dissatisfaction with the answer our normal concepts would give. What this shows is that while we may have criteria fit to judge familiar things, those criteria do not exhaust those concepts. The shape of those concepts is influenced by the conditions that surround their ordinary application. When these conditions substantially change, we may no longer be satisfied by those criteria. They may no longer capture all that our concept was designed to.

How to characterise our response to this case? One can imagine conflicting answers. Some might argue that, in confronting the zombies, we are forced to alter our existing concepts of 'living' and 'dead'. But others might suppose that, on the contrary, what our reaction shows is that the concepts were already prepared for such examples. No changes had to be made. Those features that were important in our not wanting to include zombies exclusively in the category of the dead were always important, and thus always guided our use accordingly. The purpose of our concept 'dead' concerned more than simply delineating human beings whose internal organs had permanently ceased to function, it is just that – up until now – their no longer functioning had always been accompanied by these other features and so they were never made explicit. The answer regarding who is correct here may not be straightforward. Given that we do not have to make explicit the norms by which we speak, it may be indeterminate exactly what was performing the role of criteria. This will not apply to all cases of this kind. Sometimes criteria are made explicit, and, in such cases, it may be easy to identify conceptual change and evolution. But elsewhere, it can be opaque as to what (precisely) is performing the role of criteria. It is also possible that two people use concepts according to slightly different criteria, but never discover this fact until a case like this arises. These are all live options, and which is the most appropriate description can only be judged according to the details of the case.[215]

The indeterminacy just discussed is not merely internal to the second type of open texture but affects our capacity to distinguish between them. If we can be unsure as to whether a given property is necessary or sufficient, then we can be unsure as to whether our criteria appear

[214] Waismann (1965, 332) discusses a similar case regarding people whose hearts stop beating but who otherwise behave normally. He suggests we might devise two separate concepts: 'dead' and 'heart-dead' to acknowledge the strangeness of this occurrence.
[215] Disputes of this kind will be considered in more detail in the following chapters.

to give the 'wrong' answer or fail to provide an answer altogether (for that may depend on those criteria that are unclear). The distinction between the two kinds, then, is not sharp. Even on the borderline, however, the common feature remains pronounced: a new case exposes the boundaries of a concept to a challenge they do not usually face. The shared similarity between the different cases is that the boundaries which are challenged by these unanticipated cases are (or were) generally unproblematic. Their 'openness' is an allusion to their being free from regulation due to their never being exposed.

While the examples I have explored in this section have been fantastical, it is important to note that they do not have to be. For instance, in Newtonian physics, momentum was defined as 'mass times velocity' and it was thought that the same magnitude is conserved in elastic collisions. But Einstein's theory of special relativity posed a problem, for it suggests that 'rest-mass times velocity' is not (exactly) equal to that which is conserved in elastic collisions in all circumstances. Two properties that, prior to Einstein's theory, had been thought to coincide invariably and had, therefore, been bound together were forced apart. This only became necessary after we accepted that 'mass times velocity' is not strictly preserved in elastic collision. Einstein's theory contested an otherwise unchallenged conceptual boundary that had been built upon the seeming coincidence of two properties.[216] A second case is mentioned by Kuhn. He asks '[w]hat should one have said when confronted by an egg-laying creature that suckles its young? Is it a mammal or is it not?', and notes that these are circumstances in which, as Austin put it, we don't know what to say – words fail us.[217] He concludes that:

> Such circumstances, if they endure for long, call forth a locally different lexicon, one that permits an answer, but to a slightly altered question: "Yes, the creature is a mammal" (but to be a mammal is not what it was before). The new lexicon opens new possibilities, ones that could not have been stipulated by the use of the old.[218] (Kuhn 1990, 306-7)

Concepts can be constituted and shaped by various forms of indeterminacy and vagueness. To reflect on these kinds of indeterminacies is part of what it is (or can be) to answer conceptual questions. There is no inherent tension in the idea of norms as imprecise guides to a concept's application. Such norms determine the shape of these sometimes loosely defined, or open textured, concepts. Whether it is appropriate to apply a given concept still depends on those norms, but there can be occasions when the norm provides no clear guidance. Categories and

[216] See chapter seven.
[217] Kuhn (1990, 306) and Austin (1961, 56).
[218] Cf. BT 196.

classifications are often – at the margins, at least – not sharply defined. They are especially prone to being exposed by unanticipated cases, for language is developed to be used in the circumstances we find ourselves in, not the possible ones that can be described but rarely, if ever, are realised. Both vagueness and open texture, then, can come to characterise our concepts. Their presence is no more a threat to the possibility of explaining necessity via constitutive relations than anything else we've encountered. For again, the messiness of a norm's content is not a challenge to its constitutive role – rather it shapes, and is therefore indicative of, a concept with those characteristics. As we will see, some necessities may be domain-specific, or apply only to certain uses of a given word (without accounting for its entire meaning).

6. Conclusion

This chapter has sought to explore the diversity within the workings of language. It is by no means an exhaustive list of the way language may fail to conform to expectations of orderliness or rigidity. Such expectations are misplaced, and they are not helped by the notion that language is 'rule-governed'.[219] For rules are typically explicit principles by which we do something. Whereas in language there are ways of doing things correctly that are rarely set out by rules we pay attention to. Moreover, it is constantly evolving. We may project words beyond their usual surroundings without ever having to make explicit the rules by which we are doing so. Rarely will there be any. Instead, competent language-users will recognise the relevant analogies at a glance and have no need to press the speaker to define under what circumstances it would be appropriate to make the same move elsewhere.[220] There is a piecemeal quality to language that is hard to account for in the thought that it is rule-governed. If we are to talk of rules in language, it must be clear that they are often in stark contrast to the rules of calculus.[221] It must also be clear that, while there are standards, or norms, throughout language, necessary to distinguish the meaning of one word from another, it does not follow that this requires rules that define

[219] Schroeder (2013) provides reasons for caution when thinking about 'rules' of language. See also Schroeder (2017).

[220] It is sometimes argued by utilitarians that they cannot be accused of supporting the punishment of innocent people, for the term 'punishment' implies that the party is guilty – see Quinton (1954, 137). Accordingly, it is not said to be a peculiar injustice to 'punish' someone who is innocent, but instead a mere harm against someone, who has done nothing wrong, in service of the greater good. This response misunderstands the very flexibility of language I have been concerned with, and language's propensity to be stretched and projected into contexts to make important connections clear. That it is different from punishment as it is usually understood is precisely the point its critics make. The thought is that if the acts associated with punishment are otherwise reserved for the guilty, why should we be comfortable with them in the context of innocent parties? Their point is that it is for good reasons that punishment is reserved for guilty parties (there is an ethical basis for our concept, we might say) and the utilitarian fails to pay appropriate attention to those reasons.

[221] Mulhall (2001) discusses the differences between language and calculus.

necessary and sufficient conditions for their application in all circumstances. Simple rule-formulations are not always possible, and often they are mere approximations of correct word-use. And with language's propensity to evolve and its potential to be projected into new contexts, those approximations may need to be revised.

It is critical to notice just how messy the distinction between conceptual criteria and empirical symptoms can become. It is not of great surprise that some have come to doubt this distinction altogether. If language were cleaner and, for instance, governed by clear and simple rules, there might be little room for confusing criteria with symptoms. But I hope to have shown in this chapter that the messiness can be explained while maintaining the *in-principle* distinction between the two. Our systems of representation are such that they can give rise to this kind of indeterminacy, yet it can be explained by employing the distinction (between conceptual norms and empirical propositions) that might otherwise be said to be in danger of collapsing. In the next chapter, I respond to those who argue that linguistic indeterminacy *does* make the distinction between conceptual and empirical questions, and so the linguistic account of necessity, untenable.

(6) Quinean Qualms

Quine famously rebukes the distinction I have so far promoted. In 'Two Dogmas of Empiricism' and other works,[222] he questions the analytic-synthetic distinction, thereby casting doubt on the possibility of distinguishing empirical propositions from conceptual ones. He suggests that it is impossible to tell from linguistic behaviour whether a given proposition is a meaning-determining principle, or a well-entrenched empirical belief. A metaphor from his philosophy comes in the shape of a spider's web of beliefs. Those closest to the centre are most entrenched and difficult to dislodge. Should they be replaced, this will bring into question many of the beliefs that lie around the edges. That we have well-entrenched beliefs that can only be doubted along with many others is not in question. The contentious aspect of Quine's view, so far as I am concerned, is that he thinks all propositions fit into the web. Thus, there is no distinction in kind between conceptual and empirical propositions.[223] 'Conceptual propositions' are merely some of our most entrenched beliefs.

For the sake of clarity, let us stipulate that by 'analyticity' I mean its epistemological conception. A statement is epistemologically analytic if by denying it one thereby manifests a degree of misunderstanding or employs the terms involved in a different sense. Talk of 'knowledge' and 'understanding' may, therefore, feature in what follows to an extent they have not elsewhere.[224] As stated previously, my explanation for epistemological analyticities is that the relevant propositions express norms that (partly) determine the meanings involved.[225]

Quine's attacks on the analytic-synthetic distinction, and the accounts that rely on it, are not compelling.[226] Those who believe in what Quine calls 'the myth of a museum'[227] – and the misconception that a conceptual proposition is true because of meanings which exist independently of, as opposed to being constituted by, the norm expressed (or reflected)[228] – are

[222] Quine (1936 [1966]; 1951; 1960a; 1960b). See Lugg (2012) for a sympathetic reading of the former.
[223] Cf. Quine (1951, 43).
[224] For extended discussion and debate of the linguistic account of *a priori* knowledge, see Boghossian and Williamson (2020).
[225] Cf. Glock (2003a, 84).
[226] Grice and Strawson (1956) and Glock (2003a) both present convincing arguments for the failure of Quine's attacks. I do not rehearse them here, though they feature in what follows. Boghossian (1996) also challenges Quine's attacks on analyticity and synonymy. Sober (2000) likewise presents a critique of those attacks and Quine's epistemological holism, which is central to one of his objections to the analytic-synthetic distinction. Moore (1997 [2019]) raises further doubts internal to Quine's project. He questions whether Quine can consistently draw the underdetermination/indeterminacy distinction without appealing to something like the analytic-synthetic divide. For some detailed criticisms of Quine's arguments about meaning, reference and translation, see Soames (1999).
[227] Quine (1968, 186).
[228] This is intended to acknowledge the point that some conceptual propositions follow from linguistic norms without themselves being properly described as norms.

certainly well challenged.[229] But I have suggested that linguistic norms, which are expressed by conceptual propositions, constitute the concepts contained within those propositions. Words do not have independently existing meanings that make us accept those norms as true, rather a community's acceptance of those norms gives those words their meaning. What follows from those norms (i.e., what is involved in accepting them) is likewise constitutive of the relevant meanings.[230]

The purpose of this chapter is not to rehearse old arguments against Quine, but rather to use his insights to pose specific challenges to my account. The general attack on the distinction between (what I call) conceptual and empirical propositions is hyperbolic. It is the distinction in particular cases that can be called into question, not the distinction in general.[231] There are many cases where it is obvious that a proposition is more akin to a definition than an empirical fact. Many have been provided throughout this thesis. Indeed, Quine himself is forced to admit that linguistic stipulations are an example of the phenomenon he is otherwise sceptical of.[232]

Furthermore, this position collapses the distinction between necessary and contingent truth: there are no (logically) necessary truths, only strong empirical generalisations. However, the distinction between necessary and contingent truth should not be collapsed but modified.[233] Quine is right that 'necessary truths' are revisable, insofar as they express linguistic norms that can be abandoned (how well-entrenched those norms are will vary). As Huw Price writes, 'linguistic rules are never absolute, and pragmatic restructuring is never entirely off the agenda.'[234] We should agree with Carnap, as Price does,[235] and hold that rules are laid down with the

---

[229] Schroeder (2009, 107) suggests that there might be a sense in which conceptual propositions are true in virtue of meaning, if we think meaning is determined by use.

[230] Like Quine, Waismann (1949, 27-8) criticises this view, as does Wittgenstein in his discussion of 'meaning-bodies' (PG 16 and PI §559). Some Wittgensteinians argue that Carnap subscribed to something of this view. Hacker (2003) challenges Conant (2001) on this point but only as it pertains to discussions of nonsense. Hacker argues that Carnap (at least in 'The Elimination of Metaphysics') does not think expressions are meaningless because their meanings do not properly fit together. Hacker (ibid, nn. 6) does, however, suggest that Carnap can be accused of the meaning-body conception in his account of analytic truths, because he (allegedly) thinks their truth follows from the meanings of words. But in *The Logical Syntax of Language*, Carnap (1934 [1937], xv) appears to relieve himself of any such notion when he writes: 'let any postulates and any rules of inference be chosen arbitrarily; then this choice, whatever it may be, will determine what meaning is to be assigned to the fundamental logical symbols.' Cf. Carnap (1934 [1937], 17) where he seems to suggest the meaning of mathematical concepts is given by rules. Glock (2003b, 156) appears to agree that attributing the meaning-body conception to the logical empiricists does not, in the end, stand up to scrutiny, even if there are passages that appear to support this reading. Such a view could perhaps be attributed to Frege (1893/1903 [2013]), though see Lavers (2019) for potential complications.

[231] Cf. Warren (2018, 439-40).

[232] Quine (1951, 26). Lycan (1994, 264), however, argues that 'a good Quinean ought to argue as powerfully as possible against the possibility of creating genuine analyticities even by explicit stipulation'. He may be right about a good Quinean, but Quine himself was not willing to make the same mistake.

[233] Cf. Warren (2022) who says something similar about the *a priori*.

[234] Price (2009, 326).

[235] Price (1997).

reservation that they might be altered in the future when it is judged expedient to do so.[236] One might, then, retain something like the analytic-synthetic distinction and instead suggest that its primary function is to tell 'us what is involved in revising a given sentence'.[237] We might, for example, say that a conceptual statement cannot fail to be true while the concepts involved remain fixed and that this is what necessary truth consists in, as I have argued.

1. Moderate indeterminacy

Quine's philosophy is vast and complex and I have no intention of discussing it in detail.[238] Instead, I want to present arguments against his conclusion that one can get by without meaning-determining principles. Central to Quine's challenge is that a study of our linguistic practice reveals that the distinction cannot be made. What Quine effectively denies is that there are any truly meaning-constituting relations that hold between principles and concepts, for there are no such principles to stand in those relations. This, I will argue, undermines the practice we are known to engage in.

Boghossian is right, in my view, that should one reject the distinction between meaning-determining principles and others, one will be committed to Quine's more extreme (and less popular) thesis of radical meaning-indeterminacy. Boghossian provides a pithy summary of his argument:

> If a subset (not necessarily proper) of accepted principles is supposed to determine meaning; and if there is no fact of the matter as to which subset that is; then there is, to that extent, no fact of the matter as to what meaning has been determined. (Boghossian 1996, 385)

Or, as Warren says, 'if you think there are facts about meaning, then there will be facts about which alterations in our attitudes force a change of meaning and which do not.'[239] And this very distinction can allow you to identify conceptual propositions.

Not everyone, however, agrees. Giaquinto argues that moderate indeterminacy suffices to undermine the distinction between meaning-determining principles and their empirical

[236] Carnap ([1934] 1937, 318).
[237] Yap (2010, 445).
[238] His views change over time, as prominent critics show (e.g., Glock 2003a and Westphal 2015). Indeed, Hylton (2019, 2) suggests that 'from 1974 on, Quine accepts a version of analyticity' (see Quine 1974). Nevertheless, I hope the objections I raise in his spirit are at least recognisably *Quinean*.
[239] Warren (2018, 440).

counterparts.[240] He claims that, while there are some semantic facts, meanings are not fine-grained enough to provide knowledge of neat, analytic truths. On the contrary, we are often unable to say whether a given statement is meaning-determining or not. Seldom is there – to put it in Giaquinto's favoured terms – a fact of the matter as to whether something can be known via our semantic understanding.

The first thing to notice here is that this is not something I need to deny. If the point is that there are some propositions for which we cannot say whether they are meaning-determining or not, then that has already been conceded. The point Boghossian would surely want to make, in accordance with the quotation above, is that the extent of said indeterminacy will be reflected in the meaning of the corresponding word(s). Suppose we cannot say whether 'All *X*s have property *Y*' is meaning-determining. This means it is indeterminate as to whether to qualify as an instance of '*X*' something must have property *Y*. If this is true, then something lacking property *Y* cannot be counted as an '*X*'. But we are not licensed to rule it out either. Assuming there are no other principles to settle the matter (suppose the phenomenon meets all other criteria), there is no determinate answer. It is a genuine borderline case. Hence, when we talk about '*X*s', it is unclear whether that includes the borderline case.

This echoes a point I have laboured to make, namely that forms of vagueness and indeterminacy can themselves shape the concepts we use.[241] It is our semantic understanding that gives rise to knowledge of these indeterminacies. Far from being problematic, they characterise the relevant meanings. In some cases, we may seek to eradicate indeterminacy; in others, we seem happy to live with it.

1.1. Giaquinto and multi-faceted meaning

But Giaquinto has another kind of indeterminacy of meaning in mind, which much of his paper focuses on. He opens his paper with a comparison between 'cow' and the French word 'vache'. He suggests that one might initially think they share the same meaning, namely 'female bovine'.

---

[240] Giaquinto (2008).

[241] There is a question here about how to distinguish imprecision, vagueness, and indeterminacy. The latter is merely the notion that there is no fact of the matter (whatever the matter be). Vagueness, in this context, comes in degrees and is a property of concepts that have imprecise (inexact) boundaries. Those boundaries can be represented by 'meaning-determining principles' and are defined by whatever norms we use the concept in accordance with. These standards can vary in vagueness and precision. The vaguer the standards, the more room is left for indeterminate results when judging whether it is appropriate to apply the corresponding concept to a given case. Another consequence of vague standards can be indeterminacy with respect to whether some given principle is meaning-determining or not.

However, it is explained that 'cow' is used in several ways 'vache' is not. For example, 'cow' can also be used as a verb ('to be cowed into silence'). Giaquinto is happy to park this issue because he thinks it is obvious that we would say these are two different meanings. The reason being that they belong to distinct syntactic categories and do not seem to share etymological roots. It seems to be a quintessential case of homonymy.

What is of more interest for Giaquinto's purposes is that 'cow' is also used in English to talk about animals other than bovines, including female elephants and whales. In this case, it is far from obvious, he claims, that we would be inclined to talk about several different meanings. Whether uses form a single meaning is determined, he tells us, by their 'conceptual connection[s]'.[242] His argument in this case is not that we would count only one, but rather that we would not know what to say. The reason for this is that there is no fact of the matter as to whether these uses are aspects of a single meaning. We lack precise enough criteria to individuate meanings in this case (and many others). Hence, we have no principled way of deciding whether 'cow' is polysemous. For this reason, he argues, knowing the meaning of 'all cows are female bovines' is not sufficient for knowledge of the statement's truth. The implication is that this lesson would apply across much of our language.

This is a curious argument. For one, all of this can be granted without any damage being done to Giaquinto's targets. One can hardly charge Boghossian on the grounds that we cannot know via our semantic understanding that all cows are female bovines if, in fact, all cows are not female bovines. One cannot be expected to provide an account of our knowledge of something that isn't true.

This response is given by Nyseth, and while successful on its own terms, it may concede too much.[243] For where 'cow' is used according to the relevant restrictions, is it not true that all cows are female bovines? Insofar as we are using 'cow' according to those restrictions (to talk about *those* creatures), it seems to follow trivially. Perhaps one is worried here about the move from 'meaning' to 'use', but the origin of that shift lies in Giaquinto's analysis. It is through understanding what the sentence 'all cows are female bovines' *says* that we are supposed to know it is true. But, according to Giaquinto, the meaning of 'cow' is a wider summary of what it says in different circumstances. Now, if we recast 'meaning' in such a way that it cannot stand for what any given word or sentence says in a particular context, then of course we are going to have to revert to what Giaquinto calls 'uses'. But it is hard to see why one should feel challenged by what

[242] Giaquinto (2008, 94).
[243] Nyseth (2017).

is – in effect – a terminological switch. So, while an argument for 'all cows are female bovines', according to Giaquinto, cannot be made like this (because the second premise is false):

> All cows are cows.
> 'Cow' means the same as 'female bovine'.
> So, all cows are female bovines.

We could form a sound argument that stated how 'cow' was being used. Moreover, it is our semantic understanding (which grants us access to complex meanings allegedly composed of many incompatible uses) that accounts for our knowledge of its conclusion. Given that word-meaning is said to be composed of the different uses (each constituted by a set of use-determining principles), it seems appropriate to say each use constitutes some aspect of the word's meaning. In that respect, use-determining principles can themselves be regarded as meaning-determining principles. Another way of seeing this is to recognise that a meaning can, in theory, be constituted by a single use. In that case, there would be just one set of meaning-determining principles, namely those that govern the word's use.

The possibility Giaquinto alludes to is that within single meanings, several incompatible uses may be present. His claim is that, in some cases, there is no fact of the matter as to whether there are several meanings or just one. But all this amounts to is an increase in the complexity of the relevant semantic facts. Moreover, we still need a way of expressing those principles that determine different uses. There is a use of 'cow' such that it only applies to female bovines. The indeterminacy regards the question of whether other uses of 'cow' can be counted as forming part of the same meaning. Hence, the principle that determines the individual use ('female bovine') still determines, by Giaquinto's own lights, one aspect of the relevant meaning(s). Indeed, his discussion of meaning as an input-output function suggests as much. The thought is that meaning is a function that includes many conditionals of the form: if 'utterance context and pragmatically relevant background beliefs' are such and such, then use is governed like this.[244] But this poses no serious threat to the distinction between principles. It merely challenges the idea that meanings can be individuated by the norms that determine what is said when a given word is used. If Giaquinto is right, a full specification of meaning would presumably require a survey of all uses and the different restrictions that apply to them. It would also demand an acknowledgement that there is no fact of the matter as to whether there is a single meaning or

[244] Giaquinto (2008, 94). The alternative he works with appears to be that meanings are simply composed of several different use-principles. I don't take these suggestions to be in conflict.

many. But this hardly stops those individual restrictions from having a meaning-determining role. Rather, they would characterise one aspect of a word's complex, multi-faceted meaning.

1.2 Against Giaquinto's revision

Whether it is advisable to accept Giaquinto's way of presenting things is another matter. He argues that 'cow' doesn't mean 'female bovine' because there are other, related uses of 'cow' that refer to different animals. There is, however, an alternative way of responding to this fact. We might say that our meaning explanations identify *a* use of the word, allowing that the same word, or at least the same sequence of letters, may have other uses that are irrelevant in this context.

One reason this might be preferred is that we can ask for the meaning of 'cow' inside a particular context and expect the answer not to provide a function that has several outputs, one of which is 'female bovine'. For we will be asking what the meaning of this particular use is (because we want to know what is being said). On Giaquinto's account, knowing what is said in a specific context requires us to know how the term is being used, but this is separate from word-meaning. Meaning is the function from which uses are determined.

It might be thought that indexicals assist Giaquinto's case here. After all, statements containing indexicals may have the same meaning yet express different propositions. Hence, knowing the meaning of 'I', for instance, doesn't necessarily tell us what is being said (we need to know who is speaking). But these cases are not obviously analogous. One flawed argument would be to say indexicals have a single meaning that can be given by a simple principle, whereas Giaquinto's principles governing 'cow' do themselves change from context to context. It is flawed because, on Giaquinto's understanding, it isn't so much that the principles change, only that they are complex enough that in different contexts different restrictions apply. The principles may not be simple, but so what? The thought is that an *unchanging* function determines different parameters for different uses.

This (correct) response, however, exposes the real problem with the analogy. In the case of indexicals, you can use 'I' as I do and say something different. If we both say 'I have brown hair', for example, we have used the words in the same way but expressed different propositions. This is clearly not the case for Giaquinto's 'cow'. You couldn't take my use of 'cow' as in 'female bovine' and use it to refer to something else. In other words, there is a distinction between different uses and different people employing the same use. This need not be news to Giaquinto, but it does mean his revision cannot be made more intuitive by this comparison with indexicals.

In any case, there is a compelling reason to reject Giaquinto's account. He says that we 'cannot reasonably claim that there are several homonymous nouns "cow", one for each of the relevant species'.[245] Now, paradigmatic cases of homonymy are indeed, as Giaquinto says, ones in which the meanings share very little relation (as in baseball *bat* and fruit *bat*). But that the two related uses of 'cow' are not homonymous could only tell against my interpretation if homonymy were merely a matter of the same sequence of letters having *any* difference in meaning. Yet if this is all homonymy comes to, then those uses are indeed homonyms. It is indisputable that there are distinct meanings (and concepts) relating to our different uses of 'cow'. We can use 'cow' so that it is restricted to the very same standards we employ when saying 'female bovine'. We can also use it to standards that incorporate a wider group of animals. No one would suggest that 'female bovine' means the same as a word that includes members of that wider group. Yet 'cow' is used in some contexts in precisely the fashion that implies it means the same as 'female bovine'. So, there is no question that different meanings relate to our different uses. We already accept there are distinct meanings in play for these different patterns of use. The controversy concerns only what we say, in the abstract, about a word that can be used both ways. This may be an interesting question for lexicographers, but its philosophical import is minimal.

1.3 A Waismann-inspired alternative

The kinds of observations that lead Giaquinto to his conclusion are not without precedent. Moreover, there is precedent within the tradition he targets. Waismann, author of *The Principles of Linguistic Philosophy*, makes precisely these types of observations in a series of papers on the analytic-synthetic distinction, published in *Analysis* between 1949 and 1953.[246]

One interpretation of what is going on in Waismann's papers is presented by Gillian Russell.[247] She suggests we understand Waismann as arguing that some words with many, incompatible uses have 'incomplete meanings'. To deploy such a word in a specific context is a matter of 'precisifying' that incomplete meaning. Thus, she provides a model for how we might try to understand the possibility that words like 'cow' have a single, multi-faceted meaning, which allows the term to refer to the wider group of animals in some contexts while restricting it to

[245] Giaquinto (2008, 92).
[246] Waismann (1949; 1950; 1951a; 1951b; 1952; 1953).
[247] Russell (2019).

female bovines in others.[248] Following Russell's Waismann, we might say that we can choose to 'precisify' its meaning in one of two incompatible ways. That all cows are female bovines, we might say, is necessarily true on the latter specification, but contingent (and false) on the former. Thus, following Russell's Waismann, whether the statement is analytic or synthetic will depend on how we choose to 'precisify' the meaning. According to Russell, 'though Waismann would claim that we haven't *changed* the meaning of the words in the sentence (because the incompleteness of the meaning means that we didn't need to do anything that *conflicts* with the normal usage patterns) we have still chosen to precisify the incomplete meaning in one of two incompatible directions.'[249]

Notice that Russell's Waismann arguably avoids the pitfalls of Giaquinto's analysis by claiming a word's *meaning* is 'precisified' in different directions. It also has the advantage of clarifying the position of principles that do determine meaning in particular contexts, by granting that they are themselves to be considered analytic. But while the account presented by Russell is an improvement, and the phenomenon she identifies real, it still leaves something to be desired. Firstly, we are owed an account of 'complete meaning' for the distinction to be illuminating. Why couldn't multi-faceted meanings be complete? Secondly, incompleteness is not obviously remedied by additional precision, unless what is meant by 'incompleteness' is 'imprecision'. However, we then meet a third worry: it is unclear that the process of determining a word's meaning in a particular context is a matter of making something more precise. A multi-faceted meaning may be perfectly precise but license several, incompatible (yet precisely defined) uses. It is, therefore, better described as a process of determining, or specifying, what the word is to mean in a given context, acknowledging that the word is already prepared to mean several different (but related) things depending on the circumstances.

1.4 Diagnosis and evaluation

Giaquinto is exercised by the problem of counting meanings as though trying to compile a dictionary, ignoring what we say when words are restricted within particular contexts. It might be that a dictionary would count only one meaning of 'cow', covering incompatible uses of that term (ignoring the verb). The dictionary gives an overview of the term's multi-faceted meaning.

[248] To be clear, I am not arguing there is, in fact, only one meaning here. I am considering a framework Waismann provides for understanding the possibility of a single word-meaning licensing several different, incompatible uses of that same word.
[249] Russell (2019, 186). For the avoidance of doubt, I am not claiming that Russell or Waismann would accept Giaquinto's example as falling into this category.

But is there any reason to give what we'd say in those circumstances priority over what we say inside particular contexts? The process of counting meanings in a dictionary and applying one's findings to individual sentences only succeeds if such a reason can be found. Moreover, the Waismann-inspired alternative suggests that giving a single, overarching meaning for a term in the abstract doesn't exclude the possibility that the term might be made to mean different things in specific circumstances.

There are indeed words with established uses, as Waismann appears to suggest, that genuinely hover between different concepts.[250] In context, it is determined which of the concepts is being deployed. Consider Giaquinto's example of 'book': sometimes we refer to the abstract text one is said to author, other times the physical copies one owns. In some contexts, we may refer to both: 'the book was brilliant, but the font horrible!'. No one can doubt that the meanings of 'abstract text' and 'physical copy' are distinct. The point of contention is what we should say about the meaning of 'book'. Russell's interpretation of Waismann, augmented in the ways I've suggested, remains the more compelling because it accommodates what we say both in and outside of context. Moreover, it acknowledges that single uses that compose a broader, multi-faceted meaning can themselves be identified with meanings that we already recognise as distinct ('abstract text' versus 'physical copy'). An explanation of multi-faceted meaning thus comprises a statement of what the word can be used to mean in different contexts.

In the end, then, what Giaquinto's discussion reveals is that we don't have a precise method for individuating meanings for words outside of context. This is perhaps unsurprising. A convention that determined as much would not be of any great value. Inside contexts, it is different. Knowing that someone means something different by their use of the same word is important for understanding what they have said. This is true even if the difference in meaning is small.

1.5 Conclusion on moderate indeterminacy

Giaquinto's strategy against the linguistic account is to show how few neat and precise linguistic rules there are. The thought is that very rarely can we say something simple like "'*A*" means the same as "*B*", so all *A*s are *B*s'. But the linguistic account does not require such simplicity. It

[250] Indeed, many words seem capable of having context-specific restrictions placed on them that don't generally apply. These different senses won't necessarily be picked up by dictionary definitions – they may even be particular to individual conversations. We do, after all, sometimes stipulate that by some common word we will mean something more specific than usual.

merely requires norms that determine what '*A*' means. From those norms, it is possible to derive what can and cannot count as '*A*', as well as, perhaps, things for which it is unclear either way. Nothing Giaquinto says gives us any cause to doubt that knowledge of these possibilities arises from our semantic understanding. It might help to distinguish two questions that are relevant here:

(i) what is the nature of language, linguistic meaning, and linguistic norms?

(ii) what accounts for *a priori* knowledge (or truth)?

It is clear how a form of Quinean scepticism about meaning in response to (i) could affect our answer to (ii) (because meanings could not play any role). In contrast, Giaquinto fails to show how his answer to (i) affects our answer to (ii). There is no compulsion for linguistic philosophers to see language as a neat, regimented calculus. On the contrary, several philosophers belonging to this tradition have been the most powerful critics of this picture of language. Among those we can include Waismann and Wittgenstein (in his later work). Here's Waismann:

> It now begins to dawn on us that we are confronting one of the big temptations which has misled so many philosophers in the past: namely the craving for a fixed formula comprising, as in a magic crystal, the whole, often so infinitely complicated and elusive, meaning of a word. (Waismann 1965, 85)

In his excoriating review of David Pole's *The Later Philosophy of Wittgenstein*, Stanley Cavell laments that the presentation of Wittgenstein is 'ironically blind'.[251] The chief complaint Cavell has against Pole's reading is that it depicts Wittgenstein's conception of language as being governed by a complete set of rules, where it is always obvious when a rule has been followed. Far from games posing as helpful comparisons, language morphs into a game of the most regimented form. On this reading, the kinds of indeterminacies we have considered so far in this thesis leave no trace in Wittgenstein's thought. Cavell draws the very opposite conclusion:

> That everyday language does not, in fact or in essence, depend upon such a structure and conception of rules, and yet that the absence of such a structure in no way impairs its functioning, is what the picture of language drawn in the later philosophy is about. […]

[251] Cavell (1962 [2002], 44).

> Whether the later Wittgenstein describes language as being roughly like a calculus with fixed rules working in that way is not a question which can seriously be discussed.[252] (Cavell 1962 [2002], 45)

Paul Livingston follows Cavell in asking similar questions of William Child, in a review of the latter's *Wittgenstein*. According to Livingston, Pole's and Child's presentations of Wittgenstein suggest 'it is readily possible to determine the correctness of performance by reference to well-defined standards or rules'.[253] Livingston, in order to reclaim Wittgenstein from this misinterpretation, contrasts him with Carnap, who – it is alleged – believes us capable of easily distinguishing between questions he counts as 'internal' and 'external'.[254] It is debatable whether this is even a fair reading of Carnap, let alone Wittgenstein.[255] There are countless passages where the latter disavows himself of any such conception of language.

> For remember that in general we don't use language according to strict rules – it hasn't been taught us by means of strict rules, either. (*BB*, 25)

> When we talk of language as a symbolism used in exact calculus, that which is in our mind can be found in the sciences and in mathematics. Our ordinary use of language conforms to this standard of exactness only in rare cases. (*Ibid.*)

> … in philosophy we often *compare* the use of words with games, calculi with fixed rules, but cannot say that someone who is using language *must* be playing such a game. (*PI* §81)

> I do not want to reduce unsharpness to sharpness; but to capture unsharpness conceptually (MS 137, 64; cf. *ibid*, 105; PI §76).

2 Radical indeterminacy

[252] Cf. (*Ibid,* 48). The first part of Mulhall (2001) explores, in conversation with Wittgenstein, the differences between calculus and language.
[253] Livingston (2011).
[254] This is not an endorsement of the mistake Bird (1995) attributes to Quine of thinking the analytic-synthetic distinction is equivalent to Carnap's internal-external distinction. (He impugns Hookway (1988) for making the same mistake.) The crux of Bird's argument is that, on Carnap's account, some internal questions will have analytic answers and others will have synthetic ones. For instance, in formal systems of logic and mathematics, the answers will be analytic. But in a thing-language, even general internal questions, such as whether there are chairs, are synthetic: their answers will depend on what exists. The answer will be analytically derivable from particular ones, but that does not make them analytic. Thus, internal answers can be synthetic or analytic, and any attempt to lean on the notion of 'generality' to equate the two sets of distinctions will also fail as answers to general internal questions can be analytic or synthetic too. Carnap, so Bird argues, uses the analytic-synthetic distinction primarily to contrast formal languages with empirical ones, whereas the internal-external distinction applies within each single language.
[255] Bird (1995, 59) suggests it might not be a fair reflection of Carnap's view.

What of the more radical Quinean critique? The challenges to my account that I take to be in Quine's spirit are simple: if we accept that there is no determinate answer as to whether a given feature is a criterion or symptom in certain cases, how can we (i) nevertheless apply the distinction in those difficult cases and (ii) be sure that, in cases where there seems to be a determinate answer, there really is? Once we grant that it is unnecessary for there to be a definitive answer, might that not begin to undermine our confidence in the other cases? Would we not be just as prudent to treat those as similarly undetermined, where we cannot be sure that a seeming criterion is not really a very (perhaps invariably) common symptom?

Combined, these questions make for a formidable opponent. For if the distinction is irrelevant in difficult cases, and the difficult cases cannot be adequately distinguished from the easy ones, then it may turn out the distinction is irrelevant in *all* cases. We would have no distinction, then, between meaning-determining principles and others. In response to these questions, I will argue not only is it possible to apply the distinction in those difficult cases but that it is *necessary* to. That is, accurate descriptions of our linguistic practice require the distinction. In chapter eight, I discuss some ways of identifying conceptual criteria, as a way of assuaging worries expressed by the second question.

2.1 The indispensable conceptual role

The combined objection is unconvincing. It is true, as I have acknowledged already, that there are circumstances where one cannot say of some feature whether it is a criterion or symptom, or whether some statement is meaning-determining or not. But that there can be *some* such instances does not mean that that is all there is. Even in cases where we have only a loose set of features that are typically characteristic of some phenomenon, but no clear indication of which play the role of criteria, there is still a meaning-determining principle to this effect. While this leaves much open, it also rules out a great number of features that do not appear on the list. Moreover, it can still be clear that a given case does not count as an instance of that phenomenon (because it has none of those features), or that it does (because it has all of them). That is, one can have definitive verdicts without it being obvious *precisely* what is required. This situation is likely to arise when a set of features are generally found together and so we rarely have decisions to make where only a subset of them are instantiated. That there is no specific something playing the role of criteria, then, does not mean that there is *nothing* playing a criterial role. The point of the distinction was not to say that a sharp distinction will always manifest in

practice, but rather that there is a conceptual role to be fulfilled – a regularity which determines what kind of phenomenon something of a certain type is.

The arguments that follow are designed to show that the distinction between meaning-determining principles (or something like them) and their empirical counterparts is essential for an accurate characterisation of language. I will argue that language requires such a distinction and the constitutive relations the former create. This does not mean we must refer to those principles, only that there are principles (or something like them) playing the role. Where it is indeterminate which principles are meaning-determining, meaning itself will mirror that indeterminacy. Vague boundaries are fine, and a precise set of criteria unnecessary. But not having a precise set of criteria is not like having none. Merely having a set of accepted examples, with the condition that similar phenomena also count as part of the set, is something rather than nothing.

2.2 The importance of constitutive relations

Constitutive relations, those determined by conceptual norms, are vital to language. They are not a dispensable part as the first objection would suggest. Not something that applies in some uses of language and not in others. What differs is the content and nature of those norms that make language possible. But one cannot have a language without something in that role (which is to say something would not be a language without it). This point is made powerfully by Glock, who considers a practice with only contingent connections.

> In response to utterances like 'I just met a married bachelor' or 'π is a bachelor' I might be surprised, but I could not reject it as unintelligible or demand an explanation. But a predicate like 'bachelor' is meaningful only in so far as its application is tied to that of some predicates and incompatible with that of others. Without such special connections, we could employ words any way we please, and the difference between meaningful and nonsensical uses would vanish. Under such circumstances, linguistic utterances could lead one to form expectations concerning the future behaviour of the utterer, and one could use words with the intention to cause a certain result. But linguistic utterances would merely be empirical indicators of other phenomena, just as clouds indicate rain. They would have indicative value (natural meaning), but no genuinely linguistic meaning. (Glock 2003a, 95)

Such a system is by no means impossible, but it doesn't faithfully represent what we call 'language' and cannot account for linguistic meaning. Without constitutive relations, there is no distinction between correct and incorrect uses of language. Anything goes. Yet in our linguistic activities we are licensed to apply concepts in some circumstances and not others. That something *follows* from what we say enables effective communication. And this feature of language is explicit in our teaching: we tell the next generation what is right and wrong. We insist on linguistic norms. This normativity provides a powerful tool and tearing from it the conceptual connections that make it possible leaves only a hollow shell. To be sure, we may each have our own idiosyncratic uses of language. We may apply words to slightly different standards. In general, however, these tend to be marginal, not least because such idiosyncrasies – if left to run rampant – detract from the powers that language provides.[256] We would be forever having to clarify ourselves.[257] More importantly, such idiosyncrasies still have application conditions that can be made intelligible to one's interlocutor. Those conditions set out what is and is not compatible with the use of the relevant terms, thereby explaining what those terms say. What distinguishes the mere making of noises from the speaking of words is linguistic meaning, which relies on connections stronger than mere empirical associations.[258] Linguistic artefacts can stand on their own precisely because they have a normative architecture to secure them. Language is not constituted merely by our tendencies to make certain sounds and scribbles.

Constitutive relations are characteristic of language. Relations expressed in modal terms are crucial for the very empirical propositions that we are said by some to be limited to. Robert Brandom makes essentially the same point when he says:

> The ability to use ordinary empirical descriptive terms such as 'green', 'rigid', and 'mass' already presupposes grasp of the kinds of properties and relations made explicit by modal vocabulary. (Brandom 2008, 97-8)

It is crucial to the determinate application of an empirical concept that we can distinguish between circumstances that warrant the concept's application and those that do not. This is impossible without the application of a concept carrying commitments of the kind we are

[256] Cf. Railton (2000).
[257] It is worth recognising that individually we can use different standards depending on the context we find ourselves in. We may speak according to some colloquialism in our hometown yet write our essays according to more formal, established or widely recognised standards. An unusual, idiosyncratic use may be employed by oneself under a different guise. See Waismann (1951b, 121-3). Likewise, we can at different moments hold some features to be criteria, and at other times symptoms. Cf. PI §354.
[258] Hence, Glock is wrong to say they could still use words, for words are meaningful sequences of letters (or meaningful sounds).

discussing – that if something is red, for example, it cannot be colourless.[259] Someone unwilling to accept this commitment fails to assert that something is red. To say something is red is to commit 'oneself to such consequences as that [it] is colored, that it is not green, and so on.'[260] Some conceptual commitments – at their most precise – can be expressed in the form of conditions of the following kind: 'if things are thus and so, then they are *x*; if not, they are not-*x*.' These commitments determine the content of our concepts, the meanings of our words. They explain what we commit ourselves to, what we are saying or thinking, when we apply the relevant concept. They explain what those linguistic artefacts represent, signify, or say.

Without such commitments, empirical propositions cannot be judged true or false. An empirical proposition can be true or false only if it has a determinate content, the possibility of which depends on the distinction between correct and incorrect applications of the concepts involved. What those concepts represent when combined (i.e., under what conditions it is correct to combine them) must be determined. Empirical claims say things are one way rather than another. This is only possible if the words used represent things as such. They must discriminate. To study the conditions that warrant a concept's application is to study what its commitments are (what users of that concept commit themselves to in so using it). Brandom calls this view, that what is made explicit by modal vocabulary is implicit in the use of the relevant concepts (including when they are used outside of modal discourse), the 'Kant-Sellars thesis'.

> On this view, modal vocabulary does not just add to the use of ordinary empirical observational vocabulary a range of expressive power that is extraneous—as though one were adding, say, culinary to nautical vocabulary. Rather, the expressive job distinctive of modal vocabulary is to articulate […] essential semantic connections among empirical concepts… (Brandom 2008, 98)

The extent to which this view can be credited to Kant is questionable, though some of his basic insights point in the right direction.[261] The most elementary point to be made in this regard is that Kant does not take himself to be concerned with language at all. While he thinks that *a priori* principles (rules of representation) are critical for the possibility of empirical judgement, he also thinks those principles are immutable characteristics of the human mind we use to construct the

[259] This relates to what Ryle (1950) argues.
[260] Brandom (2008, 113).
[261] Kant's conception of the *a priori* is multi-faceted. In one sense, *a priori* principles constitute the concept of an object. In another, they are non-revisable and valid for all times (at least as far as creatures like us are concerned). The insight I refer to concerns the former. The latter, I have argued, is mistaken. Coffa (1991) suggests the uncoupling of these senses was crucial to logical empiricism's emergence from its (partly) neo-Kantian roots (see also Friedman (2001, 72) in this regard). Friedman (2000, 370) traces this uncoupling back to Reichenbach's *The Theory of Relativity and A Priori Knowledge* (1920).

phenomenal world we experience. Two philosophers who no doubt played a crucial role in the development of the more convincing (semantic) account are Carnap and Wittgenstein,[262] whose 'insights on the a priori [do perhaps] belong in the same family as Kant's'.[263] Unlike Kant, both take those *a priori* principles to be rules of representation that are open to revision.[264] And neither think our systems of representation construct or impose a structure onto the world.[265]

> *A priori* statements, they asserted, do not differ from others merely in the strength with which the mind or other things entitle us or force us to believe what they say. It lies rather in the extraordinary role these statements play in the process of making it possible for us to say anything at all.[266] (Coffa 1991, 260)

Language does not get off the ground without those special relations that give language its power of expression, which requires an answer (of sorts) as to what can and cannot fall under a given concept. As we have seen, the norms that provide this power need not be neat and precise, but they must create logical relations. A contrary system implies that young children do not master a language but merely a vast array of different noises that may have associations related to them. There is limited sense to the notion that, while being faithful to this account, there exists a set of associations that can be correctly attributed to a particular 'word'. For there is no room for 'correctness' beyond the acknowledgement that a given sound is often made in a certain set of circumstances.[267] One can be correct about what a given noise is often associated with, but one cannot claim that this amounts to the correct use of that sound. There is nothing to make it so. If members of these systems began to teach their children that a particular sound was only

---

[262] Chalmers (2011b) responds to Quine (1951) by developing an account grounded in Carnap where modal notions play a key role in explanations of meaning: 'The intension of a predicate may be defined as its range, which comprehends those possible kinds of objects for which the predicate holds' (Carnap 1955, 39).
[263] Coffa (1991, 263).
[264] The relation between Kant and Carnap (and the logical empiricists more generally) is explored in Coffa (1991) and Friedman (2001). The relation between Kant and Wittgenstein is thoughtfully discussed in Glock (1997), Cavell (2002), Dilman (2002) and Hacker (1997; 2013). One topic of scholarly debate has been the extent to which Wittgenstein's later work can be understood as a descendant of Kant's transcendental idealism. This will be explored in the final chapter.
[265] Topey (2019) presents a contemporary account that is like Kant's in this respect. See Strawson (1966) for a critique of the transcendental metaphysics involved in Kant's account.
[266] It is, of course, not so much the statements that play the role but the conceptual norms that are stated.
[267] That there is no distinction in Quine between ostensive training and ostensive definition is a corollary of this point. Quine takes ostensive explanations to be a form of conditioning, and there is no room for ostensive explanations to play a normative role. This is in stark contrast to Wittgenstein, who accepts Quinean conditioning is an aspect of formative language-learning but also recognises that ostensive explanations continue to have a normative role throughout linguistic practice (when I check what shade of red your wall is by using a sample, for example). Ostension can be a matter of causal training *a la* Quine, but there is a separate activity of ostensive explanation (or definition) whereby the samples used may have continual normative significance.

appropriate in certain circumstances, this would be the first step towards a language with normative character.[268]

2.3 Vicious regress

A system of contingent associations, then, does not amount to a language. Signs and sounds may be associated with certain things – they may be indicative of, for example, our feelings. But the signs and sounds are not thereby given meaning. They do not carry content without a normative architecture around them to secure it. There is no convention, or standard, to determine when it is correct to use them, which would allow those signs and sounds to stand for something. It is one thing for a sound to usually be produced in certain circumstances, quite another for someone to be able to *say* something with that sound. This distinction is observed most easily when what one says is different from what is thereby indicated. Consider, for example, someone who has a habit of saying "I'm sure it'll be fine" when stressed. Clearly what is said here is not that they are stressed, yet it might be a very good indicator of their being so.

The foregoing is related to a regress that Crispin Wright and Bob Hale find in Quine's philosophy.[269] They target Quine's thoughts that (i) all principles and beliefs are revisable and (ii) when experience conflicts with our beliefs, we have a pragmatic choice between different ways of resolving the conflict (for our choice is under-determined by the evidence). The problem for Quine, they argue, is that one can hold our logic, general theories and empirical beliefs fixed, and instead revise what follows from those theories and beliefs according to that logic. And, as they point out, this means we cannot evaluate different choices when experience conflicts with our beliefs because we cannot determine what the consequences of making different revisions are, for what those consequences are is also a matter for pragmatic consideration. To decide on what those consequences are, we must consider the logical consequences of further alternatives, but if the consequences of those alternatives are also subject to pragmatic consideration, then we enter what looks like a vicious regress. Our decision always relies on the consideration of a further set of consequences that can only be pragmatically decided upon. There is never any firm ground for us to base our considerations on. For Quine's account to work, these questions must be taken off the table. Thus, Hale's conclusion that 'statements asserting that certain propositions

[268] No special significance should be ascribed to the process of language-acquisition for our purposes. 'What matters', as Wittgenstein (LWL, 38; cf. BB, 12-14) says, 'is what is given in the explanation'.
[269] Wright (1986) and Hale (2013, 54-59). See also Shapiro (2000, 338-9).

are logical consequences, by some specified logic, of certain other propositions' must not be revisable.[270]

While this argument may seem divorced from what I say above, it can be understood as an application of it inside Quine's theory. It is absurd to say that one might hold logic fixed and yet draw different consequences from the same propositions, for a given logic is constituted by principles that determine what the relevant logical consequences are. One would not be holding logic fixed if one drew different consequences from the same propositions.[271] And Quine's account requires this insight. It cannot get off the ground if the various options and their consequences lack constitutive principles to make them what they are. What he needs are special, non-contingent relations. He needs constitutive relations like those expressed by statements of 'necessary truths'. After acknowledging this, we can vindicate the thought I mentioned earlier: while conceptual statements can be 'revised', what is involved in their revision is quite different from revising empirical propositions. The former alters the terms in which we think (by altering the constitutive principles we accept), while the latter involves a revision of how we think the world is.

2.4 The conflation of language with theory

The Quinean view under consideration also conflates language with theory. As Donald Davidson writes, '[t]o give up the analytic-synthetic distinction as basic to the understanding of language is to give up the idea that we can clearly distinguish between theory and language'.[272] Language is, on this view, constituted by contingent hypotheses that we regard as being some of our most certain – those beliefs that are most entrenched and that we are least likely to revise.

Firstly, this seems to get norms of representation and statements typically labelled 'analytic' wrong. For they are often easily revised – more easily revised than some of our empirical beliefs. 'A bachelor is an unmarried man' is not immune from revision and the rules of obscure parts of mathematics are not central to our webs of belief. They need not play a central role in our understanding of the world to have their status.[273] The peculiar necessity of these propositions does not reside in their resistance to change, though some – such as *the principle of non-contradiction*, for example – no doubt are central to our way of thinking and are therefore

[270] Hale (2013, 60).
[271] See Warren (2018) for evidence that Quine recognises this.
[272] Davidson (1973, 9).
[273] See Sober (2000, 262-3) on this point.

resistant to change. What is distinctive about these truths is what appears to be at stake when we revise them. We can happily count 'a bachelor is an unmarried man' false, but in so doing the statement ends up saying something different. This can be explained. If we stipulate that we are holding the meaning of 'bachelor' fixed in accordance with its strict definition, it does not make sense to deny that all bachelors are unmarried men. Its denial is unintelligible because the sentence is an expression of what the word means. Something would not be a bachelor if it were married. This is in stark contrast to, say, 'the earth is round'. Here quite the opposite is true: it is very difficult to abandon (without sacrificing one's rationality), but easy to imagine false![274]

Secondly, the conflation of language and theory appears to get language fundamentally wrong. A language is constituted by a grammar and vocabulary that make theorising (and much else besides) possible. One can have many different and conflicting theories described in the same language. Furthermore, language does not predict as theories do. It provides a space of possibilities, a pool of possible descriptions, from which we can pick some that we believe best capture the world. The language itself does not privilege one set of descriptions over another as being more faithful to reality. That is something we try to do with our theories, which can only pick from the possibilities our language caters for. This is one reason why the history of science is one of conceptual innovation and evolution as well as discovery. We can, and do, develop our linguistic tools as and when required. As Waismann says, 'breaking away from the norm is often the *only way* of making oneself understood'.[275]

Finally, there is no sense in which we can say that a language is 'true' or 'mostly true'. For it is *what we say* in a language that can be true or false, not the language itself.[276] Quine suggests that languages and theories are both fabrics of sentences.[277] But a language is constituted by the principles by which we make meaningful contributions, including both semantic and syntactic rules. It is not a conglomerate of assertions. And even if we were to think of language as a set of sentences, it would be quite unlike any theory. For it would have to include the negation of each of its empirical propositions. Within the space of possibilities that language provides, every true empirical proposition will be accompanied by its meaningful (but false) negation. But within

[274] Coffa (1991, 55) captures a related point when discussing the force of an insight developed by the logical empiricists: 'Many fundamental scientific principles are by no means necessarily thought – indeed, it takes great effort to develop the systems of knowledge that embody them; but their denial also seems oddly impossible – they need not be thought, but if they are thought at all, they must be thought as necessary.'
[275] Waismann (1953, 84). See also Friedman (2001, 38-9).
[276] Cf. §PI, 241.
[277] Quine (1960b, 11; 1969, 308–11).

theories, contradictory sentences are marks of incoherence. Theories assert that something is the case and so the inclusion of their sentences' negations is tantamount to saying 'p' and 'not-p'.[278]

2.5 The limits of a system of contingent associations

We can, as I have said, imagine a system of contingent associations. We can describe a practice whereby the means of communication is an assembly of noises that people tend to react to in certain ways, out of habit or expectation. They might be able to indicate their own intentions to someone just in case that person happens to associate the noises they make with the correct behaviours. Our ability to describe this system relies on our knowing what constitutes an 'empirical association' – it relies on the normative nature of language that a system of mere association lacks. And that is the point. A system of mere association can provide an indication as to what we want, can see or hear. But it cannot *describe* or *explain* those things, for it lacks the necessary components to represent the world with words. It has not the powers of convention that stipulate what I need when I cry 'help!'. Of course, that does not mean I will not receive what I need: even those who cannot understand my words may come to my rescue. What we might say, then, is that we do not understand words in systems of mere association, though we may understand each other. We may be right to expect something because of the noise someone makes, given past behaviour. But we do not understand the expression in the linguistic sense: we do not understand what is conveyed by the noise, what is meant by the noise, for there is no content to be conveyed, no meaning to be understood. A system of mere associations is not an unsophisticated, or impoverished, language, but no language at all.

Of course, one might argue that all we have is the system made possible by contingent associations. As if understanding a language were no different from interpreting a baby's response to their environment. But the burden of proof is very much on the advocate at this stage. It is certainly not something I recognise as the language we use. Moreover, it seems at least possible for us to devise a language with normative character. It would be strange had we come to use and prefer a system of contingent associations, which lacks many of the powers we – so I have argued – ordinarily rely upon.

[278] Similar points are made by Hacker (1996b, 297-8) and Glock (2007, 387).

(7) Conceptual Change

The discourse of science has proved a fertile ground for objections to the conceptual/empirical distinction. This is for two main reasons: (i) empirical laws and linguistic norms can be indistinguishable, and (ii) conceptual change is said to be problematic for our notion of scientific progress.

My answer to (i) contributes to the discussion in the next chapter, insofar as it suggests a method for identifying the conceptual role. The chapter starts by demonstrating how conceptual change can reveal the conceptual/empirical distinction in cases where empirical laws and linguistic norms may be indistinguishable. One plausible form of conceptual change involves eradicating indeterminacy and sharpening our concepts.[279] Conceptual change is often a matter of refining concepts to make them determinate with respect to newly discovered phenomena.

I then consider objections to accounts that entertain the possibility of conceptual change in science. We have already seen that supporting the conceptual/empirical distinction makes room for the revision of constitutive principles.[280] That is, when Quineans object that scientists are free to revise whatever principles they like, the point is conceded and instead we focus on what is involved in revising the different kinds of statements. While it has been said that conceptual change would undermine our notion of scientific progress,[281] I will suggest it can be understood as an aspect of it. Once we are clear about what a conceptual framework does, it is natural to understand conceptual change as a feature of (some) scientific progress.

Finally, I answer another Quinean objection which is particularly relevant in the context of my earlier claims, namely that the distinction between conceptual and empirical statements is effectively collapsed because the criteria for settling on conceptual norms in science are the same as those used when deciding between different theories.[282] I agree that the criteria may be similar but deny that the conclusion follows.

1. Indeterminacy and conceptual change

[279] Cf. MS 137, 105.
[280] Cf. Carnap (1963, 921).
[281] These anxieties are felt acutely by readers of Feyerabend (1988) and Kuhn (1962). To what extent this is fair is up for debate: cf. Kuhn (1970 [2000], 157; 1982 [2000]; Moyal-Sharrock 2017). See Bird (2007) for a wider discussion about scientific progress.
[282] Cf. Quine (1948, 16-17).

Without straightforward and explicit standards of correct use, concepts can leave a great deal of indeterminacy in their wake. This indeterminacy challenges our ability to distinguish, in practice, empirical propositions from conceptual ones. If our criteria are unclear, then likewise the propositions that articulate the relevant criteria will be. The claim is not merely that making the distinction between criteria and symptoms is difficult, but that it can be impossible. It is not necessary for the functioning of language that we determine whether a given property is itself a criterion or symptom. The successful use of language relies on a form of regularity fit to make the distinction between correct and incorrect uses meaningful, but that does not necessitate a precise set of standards by which the concept is used. This is not to say that the distinction between the *roles* of symptoms and criteria, or empirical and conceptual propositions, is unclear.

This difficulty would seem especially acute in cases where we have defeasible criteria, some number of which must be met but not all. How could we determine such features were criteria rather than empirically associated symptoms without the criteria having been explicitly laid out? Perhaps Wittgenstein's private language arguments provide an example where this is possible.[283] Our ability to identify pain expression as a defeasible criterion can be explained by the fact that a *sincere* expression of pain is an indefeasible criterion (though not a necessary condition). Cases where a criterion is relevantly related to a sufficient condition appear to offer one answer to this question then. Likewise, if a criterion is part of a set that is sufficient as a whole but not when without one of its parts, then it seems as if each feature would be a criterion. They may not each be individually necessary if there are other sufficient sets that don't include them.

I have already argued that the mere presence of indeterminacy is not an issue for my account. Individual statements' status may be undetermined, but this is not an issue, so long as we accept the role this indeterminacy has in characterising the concepts involved. While that general argument holds, it has been said that science presents particularly acute cases of it. Hacker explains why:

> … what is empirical law and what is definition is often indeterminate in scientific theories, and this indeterminacy is typically of no importance until normally concomitant phenomena are found to diverge in exceptional cases. (Hacker 1996b, 295)

[283] For discussions of Wittgenstein's private language arguments, see Hacker (2019a), McGinn (2013), Mulhall (2007), Schroeder (2006) and Stern (2011).

Relatedly, Papineau argues that theoretical definitions are often imprecise but that this does not normally matter because it does not normally produce indeterminacy of referential value.[284] The thought here is that a natural place for indeterminacy to arise is in cases where several properties seem to invariably coincide. Suppose, for example, we start by identifying an object of a certain kind by one particular property (X), but for some time it has been thought that this property is always accompanied by another one (Y). It may never be clear whether we employ property Y as a conceptual criterion. After all, if we believe that Y is always caused by X, we can use Y to identify the relevant object just as well as we can use X. The method for identifying an object of that kind, regardless of whether Y is treated as a criterion or not, is the same. Of course, there might be other ways of determining whether one of the properties is a criterion, but there does not have to be. The point is not that it is always impossible, only that it can be, and the regular coincidence of properties is a potentially fertile ground for such cases. Indeed, any time there is a law-like relationship between phenomena, there is scope for this kind of indeterminacy.[285] Such indeterminacy is generally unproblematic for theoretical terms unless we find that they lack a determinate referent in the actual world. This can occur when those normally concomitant phenomena diverge because it is unclear precisely which of those phenomena are required for something to fall under the relevant term.

One example where it is obvious that the coincidence of properties is purely an empirical matter, as opposed to there being a conceptual connection, concerns the set of creatures with a heart.[286] Supposing all such creatures also have kidneys, there would be no confusion as to what the defining feature is (in part, because it is made explicit in the naming of the set). Thus, while it may be true that those creatures with a heart are the same creatures as those that have a kidney, their respective criteria and symptoms are nonetheless clear and distinct. No one would think 'having a kidney' is a defining criterion of 'having a heart', as shown by the imaginary case of a Martian with a heart and no kidney.

A second example we saw towards the end of chapter five tells a potentially different tale. Prior to the theory of special relativity, momentum (defined as 'rest-mass multiplied by velocity') appeared to coincide with that which is conserved in elastic collisions. But after Einstein, we think that the two quantities only coincide (exactly) when an observer's velocity

[284] Papineau (1996).
[285] Putnam (1962a) includes an example of this kind in which all bachelors suffer from a kind of neurosis (namely, 'sexual frustration'). He then explores our possible reactions to discovering a married man who suffers from that same neurosis. This is not an altogether plausible example, given how we use 'bachelor' (it has nothing to do with whether they suffer from a neurosis of any kind).
[286] Quine (1951).

relative to the object is zero. Nevertheless, in many cases these two quantities do very nearly coincide (where relative velocity is a small fraction of the speed of light). Only when relative velocity is very large does considerable diversion occur. It is an open question whether Newtonian physicists conceived of momentum as needing to fit both criteria, namely that quantity which is conserved in elastic collision and is calculated by mass times velocity or was defined simply as 'mass times velocity'. And not much hangs on it. For the Newtonian concept of momentum was bound up with an expectation that they would coincide, whether momentum had come to be defined by the coincidence of the two or not. Whatever the facts regarding how 'momentum' was defined, if indeed there are any definitive such facts, the root of the problem is the same: Newtonian physics supposed that these two magnitudes would invariably coincide. Theories (and later, discoveries) to the contrary led to a change in empirical beliefs and a redefining of 'momentum'. It now had to be calculated in recognition of special relativity, hence a relativised concept of momentum was developed.

If we consider two possible concepts of Newtonian momentum, we can see why – in practice – it does not matter how they were defined. On the first possibility, 'momentum' was defined as 'rest-mass times velocity', and its being equivalent to that which is conserved in elastic collision was an empirical hypothesis. In that case, Einstein is supposed to have shown that hypothesis is false. This concept 'momentum' thereby lost some of its application, for it was not momentum (relevantly defined) that was conserved in elastic collisions. The second possibility takes both as conceptual statements, and thus Einstein's theory limited the concept's (strictly) correct use to those cases where an object's relative velocity is zero. To be clear, the difference between these concepts is that it is correct to use the latter only where these magnitudes coincide. In cases where they do not coincide, the former concept says 'momentum' refers to whatever is calculated by 'rest-mass times velocity', while remaining faithful to the latter requires one to say there is no momentum in such cases. This would have meant not only that they had a narrower concept than they thought, but also that it had been misapplied on numerous occasions when relative velocity was not zero. For it had been thought that the two invariably coincide, and hence thought that one could arrive at the quantity conserved in elastic collision by multiplying rest-mass and velocity together in all cases, something Einstein's theory suggests is false. The root of the problem, as has been said above, is that it was thought that these two magnitudes invariably coincide. This remains an empirical hypothesis, even when the concept 'momentum' is defined as being the combination of the two. Thus, whether momentum is conserved in elastic collision by way of definition or not, both versions of Newtonian physics share the hypothesis that the two invariably coincide. We can treat this hypothesis as a garden variety empirical

falsehood within Newtonian physics, or as an empirical falsehood that also formed the basis, or rationale, for a concept that was central to that physics.

Or we can avoid this dichotomy altogether. We can accept that this particular conceptual boundary was, before Einstein, left unchallenged and therefore a decision did not have to be made with respect to it. With Einstein's theory came a relativised conception of momentum that was different from that found in Newtonian physics, a consequence of our thinking that these two magnitudes do not invariably coincide. Einstein's 'momentum' plays a similar role to Newton's concept, and shares a significant feature, namely its being the quantity conserved in elastic collisions. But it remains importantly different if we accept that 'rest-mass times velocity' played a conceptual role in Newtonian physics.

One might understand this interpretation of the evolution of the concept 'momentum' as suggesting that the Newtonian concept effectively presupposes a falsehood: it is used in such a way that assumes the two magnitudes are always equal. Given that is false, some might argue the term fails to refer altogether.[287] If the presupposition is false, they might think the term is rendered vague to the point that statements containing it are neither true nor false.[288] The concept stands in urgent need of clarification. This does, however, seem hyperbolic, given the Newtonian presupposition is (broadly) correct (and the concept, for the most part, applicable) when relative velocities are a small fraction of the speed of light. Indeed, it is thought to be exactly right when relative velocity is zero. We might suggest instead that only in circumstances where relative velocities are a significant fraction of the speed of light is it unclear what the Newtonian term is supposed to refer to. A stricter interpretation might limit this to circumstances where the two magnitudes coincide perfectly, but I am unconvinced that this is mandatory. Rather, one can say it fails to refer as precisely as we would like it to in cases where relative velocities are greater than zero. Part of Einstein's achievement, then, was to show how terms could be refined to make their use more precise and widely applicable.

2. Putnam's objections

[287] One might see some parallels with Strawson (1950) on presupposition here. Strawson's original paper formed a critique of Russell's theory of definite descriptions (1905). Further discussion and development of Strawson's account can be found in Strawson (1952), Sellars (1954) and Strawson's (1954) reply to Sellars. Soames (1989) provides an overview of presupposition.

[288] LaPorte (2003, 131) suggests something like this in the context of 'species' before Charles Darwin's evolutionary theory. Kuhn (1983 [2000], 212) likewise argues that, if Newton's second law is false, it is shown that the Newtonian terms in its statement fail to refer. Friedman (2001, 74) makes a similar claim with respect to the laws of motion and the concept 'absolute acceleration'.

Putnam employs the same example to argue the opposite.[289] One of his complaints relates to his semantic externalism. He claims that, far from meaning different things by 'momentum', Einstein and Newton are 'talking about the same good old momentum – the magnitude that is conserved in elastic collisions'.[290] Glock rightly criticises this response.

> Putnam's argument trades on the possibility of *oscillating* between two different definitions of momentum. What we are still talking about is momentum in one of the two senses that the term previously had, namely 'whatever quantity is preserved in elastic collision', while giving up the other of 'mass times velocity'. (Glock 2003a, 90)

Insofar as it is correct to suppose 'rest-mass times velocity' was a definition of momentum in Newtonian physics, Putnam's claim cannot be right. Following Einstein's theory, momentum is preserved in elastic collisions, which means it is only equal to rest-mass times velocity in certain cases. Only in those cases does 'momentum' refer to the magnitude calculated by multiplying rest-mass to velocity.[291]

One of the points Putnam wishes to make is that the notion of conceptual change is not sharp, because our notion of what counts as the 'same' concept is not. Putnam introduces what he calls 'law-cluster concepts', where one 'can re-identify a given law-cluster concept in different theories if the laws that govern it have sufficient overlap'.[292] But these insights can be captured without the need for scepticism with respect to conceptual change in this case (and many others). I have already accepted the possibility that 'momentum' was not defined as sharply as it might have been and acknowledged the real phenomenon of indeterminacy in that respect. Moreover, in earlier chapters I acknowledged that the diachronic identity of concepts may not be so sensitive to change that one cannot have the same concept whenever it is altered. Context may determine the extent to which we are willing to count concepts as different or the same. Whether we emphasise the differences or similarities will depend on our purposes.[293] What matters, so far as conceptual differences are concerned, is that they entail our saying or thinking something different from before (or at least our saying or thinking in different terms, given extensions may be identical through conceptual change).

---

[289] He takes the example from Quine (1976, 74-5).
[290] Putnam (1988, 11).
[291] The presupposition of this example, that 'rest-mass times velocity' was used as a definition for 'momentum', is critical here. For I am not arguing that theory change necessitates its conceptual counterpart. It is perfectly coherent to say that some scientific progress involves conceptual change without thinking all of it does.
[292] Putnam (1962a). The quotation is Shapiro's (2018, 118).
[293] Cf. Papineau (1996, 18-19) who suggests 'that sociological factors come into play when a scientific term that was hitherto thought to have a determinate reference turns out to be vague in a way that requires remedying.'

This is true even when the current variant has clearly evolved from the previous one. The point in the case of 'momentum' is that Einstein's theory prevented 'rest-mass times velocity' from playing a conceptual role. It is not my view that we must, therefore, conclude Einstein's 'momentum' is not the same concept, only that the dropping of one of its candidate definitions does alter our understanding (and use) of it. It is in an instance of *conceptual* change. This claim may come to no more than saying that there is, in Einstein's theory, a clear fact of the matter as to whether 'mass times velocity' is definitional of 'momentum', whereas in Newton's physics there is not.[294] A neat way of describing this is to say they are different concepts, but we can still recognise them as belonging to the same family, and we could, if necessary, use a different word for this grade of distinction (a different *conception* rather than concept, perhaps).[295]

It is not always clear how far Putnam is from this view.[296] For example, he asks why the statement 'momentum is conserved' should not have as great a right to be preserved as the statement 'momentum is mass times velocity', but – of course – I do not deny that it does. I only say that to acknowledge the conceptual role 'mass times velocity' had requires us to accept the implications of our abandoning it. And, as I have said above, I am happy to accept that 'there is enough continuity through change to justify' our thinking of it as the same concept 'momentum'.[297] But to understand the nature of that change one must understand the role 'mass times velocity' had. If it was considered a definition, which Putnam seems to accept, then I see little room to conclude anything other than the nature of that change was conceptual.[298] This does not entail that we move from one radically different concept to another, for we do just as well to talk of conceptual refinement (the outcome of which is a variant and clear successor of

[294] Though I suspect it comes to more in this case, as should Putnam if he accepts it is a definition of 'momentum' in Newtonian physics.
[295] Cf. Rawls (1999, 5).
[296] Shapiro (2018) attempts to grapple with the question of where Putnam, Quine and Waismann stand – relative to each other – on the question of meaning-change and analyticity. Shapiro and Roberts (2019) study Waismann's position in more detail. Russell (2019) also considers this question as it relates to Waismann, as we saw in the previous chapter. She characterises Waismann as thinking word-meaning is 'incomplete' and that we can choose to 'precisify' the meaning of a word in several incompatible directions (*ibid*, 186). As we've seen, I have misgivings about some of the technical details here. But this interpretation (or one like it) is quite consistent with what I have said, though whether it accurately captures the specific cases I am here interested in is an open question. It is consistent because while it might make one wary of talk about a change in concepts, it nonetheless holds that words are being used according to slightly different meanings. It just considers this matter to be internal to the word's overall, 'incomplete' meaning. In other words, it does not deny that this is a question of meaning rather than fact.
[297] Putnam (1988, 11).
[298] Something Khatchadourian (2007, 269), who is otherwise sympathetic to Putnam, readily accepts.

the old concept).[299] As Michael Friedman points out, 'successive paradigms emerge precisely from one another, as succeeding stages in a common tradition of cultural change.'[300]

For similar reasons, Giaquinto is wrong when he argues that we cannot identify conceptual change in cases of moderate indeterminacy.[301] He suggests that because there is no fact of the matter as to whether 'straight line' meant 'shortest path between two points', we likewise cannot say whether any conceptual change took place when this matter was clarified.[302] With the acceptance of the cosmic geometry of the general theory of relativity, he tells us, we no longer accept that between any two points there will be a straight line. The thought, presumably, is that geodesic curves (representing the shortest paths between two points in curved space) are not taken to be straight lines. When he considers the question of whether this is a change of belief or change of meaning, he suggests that '[m]oderate indeterminacy allows that there is no fact of the matter, and that that is why we have not found a principled way of deciding it.'[303] This is exactly wrong. A previously vague conceptual boundary, regarding the question of whether being the shortest path between two points is sufficient for a line to be counted as a 'straight line', is no longer vague.[304] Indeed, what follows from what he says is that we are definitive on the matter now: this is not sufficient for what we call a 'straight line'.[305]

One can see the difference by considering how these two distinct understandings of 'straight line' would be applied. In the cosmic geometry of Einstein's general theory of relativity, the shortest path between two points is not classed as a straight line. Yet it is surely open to someone employing the traditional understanding (taking it on Giaquinto's indeterminate terms) to say there is no fact of the matter as to whether the relevant line is straight. If I cannot say whether something is a 'straight line' just in case it is the shortest path between two points, then

---

[299] Moreover, elsewhere Putnam (1962 [1975], 311-12) writes: 'If the eighteenth-century chemist had insisted that there could not be, say, an acid too weak to turn litmus paper red [...] as he understands the term, then perhaps we should say that a change of meaning had occurred'. He simply denies that they would have so insisted (a historical matter which, for our purely logical purposes, we need not settle).

[300] Friedman (2001, 60; cf. 63; 100).

[301] For the avoidance of doubt, I am not denying that there may be cases where it is indeterminate as to whether any conceptual change has occurred.

[302] Giaquinto (2008, 98-99).

[303] *Ibid.*

[304] This might seem a strange thing to think. How could there fail to be a determinate answer as to whether a condition was sufficient for the use of our term? One potential way of thinking about it might be to consider the phenomenon of *open texture*. We might say that it was clearly sufficient when working inside the geometries it had been prepared for (e.g., Euclidean), but outside of them it is unclear whether it continues to be. That is, there is no answer, meaning it is indeterminate whether this continues to be a condition for its application.

[305] Donnellan (1962, 656-8) seems to make a similar mistake when he suggests that we cannot distinguish between a change in meaning and belief in certain cases, because present usage does not legislate over all hypothetical cases. When those cases arise, it is supposedly unclear which decision would preserve or change meaning. However, the point is that if the present use is not determinate with respect to those cases, then should we make a decision with respect to them to make our use determinate, we will have altered, or extended, the relevant meaning.

in cases where a line represents such a path, I am in no position to *deny* (as someone using the cosmic geometry of Einstein's theory could) that the line is straight. Of course, this is not to say I am able to affirm it either. Thus, there are some cases for which the traditional understanding may not give determinate results where the other does. We can, therefore, note that distinct standards apply to the uses of the term under consideration. Accordingly, what counts as a 'straight line' when those different standards are applied will vary.[306]

As we saw in chapter six, one point of Giaquinto's paper is that meaning-individuation is not a matter of fine-grained precision. So, some words can have multiple, incompatible uses that we would still be inclined to count as falling under a single meaning. The thought might be that, even if the new use of 'straight line' is incompatible with a previous one, they *could* still be regarded as different aspects of a single meaning (ultimately, his view – on this reading – is that there is no fact of the matter). But this is not something I *need* to deny. The point is that if one develops a new use for a particular word (even if it is not so radical that it marks itself out as a separate meaning for that term), one will have – at the very least – enriched or developed the existing meaning. Thus, acknowledging that meaning-individuation is not sharp does not help settle the question of whether this is a matter of meaning or fact. It at most suggests that: if it is a question of meaning, the result may not be a new, distinct, and separate meaning from the previous one.

With respect to the question of conceptual change in science, we are interested in the restrictions placed on a word or phrase as it appears in a particular sentence or theory and whether it is consistent with its use in a different one. If there is a determinate answer, this will tell us whether in using that word or phrase the new sentence or theory means something (slightly) different from the previous theory. Meaning-individuation in Giaquinto's sense is, therefore, of little interest in this respect.

An earlier example from Giaquinto's paper may help clarify this point. As we saw, he claims it is indeterminate whether 'cow' has several separate meanings which each refer to the different animals such female bovines, whales and elephants, or has a single meaning composed of these different uses. He does not deny that there are distinct uses with different restrictions in force. Instead, he claims that we lack criteria precise enough to determine whether those differences in use demarcate distinct meanings. Now, suppose I have used 'cow' to mean 'female bovine' for some time when teaching a class. But I come in the following week and say: 'we have found that some cows live much of their lives in Antarctic waters, and that there are large blue

[306] See the previous chapter and Nyseth (2017) for further challenges to Giaquinto's paper.

ones with hearts that beat fewer than ten times per minute'. Should my students take this as a change of meaning or belief? If Giaquinto's argument were to suggest there may be no answer, then that is either because it is ridiculous beyond words or answering a question different from the one we are interested in. In this case, my 'discoveries' wholly depend on the fact that I am now using 'cow' to refer to a wider range of animals. Insofar as the question of meaning-change is of interest to us in science, it applies to this level of understanding. That we can acknowledge the potential for incompatible uses with distinct referents inside a single meaning is relevant only to how we might come to describe this phenomenon. It does not help to answer whether this phenomenon is present in these cases.

Putnam also raises a second complaint against my perspective, which relates to his first, namely that it undermines our notion of scientific progress and the importance of scientific discoveries.[307] The result of Einstein's work is merely the attaching of an old label to a new phenomenon. This is misguided. That the magnitude calculated by multiplying rest-mass to velocity is not (always) preserved in elastic collision was an empirical hypothesis, which reflects the role that frames of reference play in Einstein's theory. New theories utilising novel conceptual resources can direct empirical inquiries towards discoveries that vindicate those theories and their conceptual frameworks.[308] Likewise, empirical discoveries can lead to the development of new categories and concepts. Moreover, the development of new concepts is far from trivial. They provide the tools for theorising, and making them more precise, for example, can make more accurate predictions possible. For instance, if predicting how substances will interact relies on an atomic understanding of those substances, accurate predictions will need to utilise concepts pertaining to that level of understanding, concepts that could not have been defined with anything like the required precision prior to significant discoveries.[309] The basis for altering Newton's concept 'momentum' resides in the fact that it is not well-suited to describing certain phenomena. As was mentioned above, the upshot of Einstein's theory was that the old concept no longer applied as neatly as was previously thought. In other cases, scientific concepts are abandoned or changed because the phenomena they aim to identify do not, in fact, exist. This does not render those concepts unintelligible, but it does make them useless in our attempts

[307] I try here to remain neutral on what such progress consists in and thus try to ensure my answer can remain consistent with more robust conceptions, such as those that take it to be a matter of acquiring cumulative knowledge. See Bird (2007) for a defence of this conception and exploration of others.
[308] Shapiro and Roberts (2019, 205) discuss this phenomenon with respect to the concept 'simultaneity' in Einstein's physics. See also Waismann (1952, 8-9), who is the subject of Shapiro and Roberts' chapter.
[309] Cf. Klagge (2017) who also suggests scientific advance can lead to modifications in our conceptual resources. An example of Wittgenstein's he cites asks whether our criteria for pain might change if we could observe the workings of nerves.

to describe the world. One potential example of this kind can be found in the very history we are studying. Einstein's theory seemed to imply that there is no such thing as the ether, meaning it (as it was then conceived) had no use inside his physics. Another example is 'phlogiston', which was a substance posited by eighteenth century scientists thought to exist in all combustible bodies.

Once a theory's terms are properly understood, there remains the crucial question of how accurately it represents the world. To what extent it coheres with observation, makes reliable predictions, and can be supported by the results of experimentation. Different theories meet these standards to differing degrees, and this is no less true when competing theories use different lexicons. Progressive, conceptual change is not supposed to replace the goal of increased accuracy, but at its best it does, perhaps, give one the resources to speak with greater accuracy. In its crudest form, this might be because older concepts were designed to refer to things that do not exist. Another possibility is that some conceptual changes provide greater precision. Of course, being in possession of more precise concepts is no guarantee that one will make accurate statements with them. The point is merely that vagueness can, in important respects, hamstring our attempts to make accurate descriptions. After all, as we saw in chapter five, if a term is sufficiently vague it can lead to statements being neither true nor false.

3. Another Quinean qualm

Penelope Maddy, channelling an argument she takes from Quine, suggests a problem for this response.[310] Recall that Putnam's problem is that conceptual change appears to rid science of all its progress. Maddy's argument suggests that my reply may save scientific progress only at the cost of dismantling the distinction I want to uphold.

> The problem for the twentieth-century Kantian is that the criteria for modifying our intuitive and categorical principles are indistinguishable from those for modifying our scientific beliefs generally. This Quinean objection to Carnap is that the criteria for adopting linguistic frameworks are indistinguishable from the criteria for adopting scientific hypotheses generally. In both cases, the cherished distinction seems groundless. (Maddy 2000, 106)

[310] See also Shapiro (2000, 343) on Quine.

While this argument may have intuitive appeal, I have provided a clear reason for the congruence of these criteria. If we are discussing scientific terms, then how useful they are will be – in part – determined by the theories they can be used to formulate. This does not collapse the distinction between the theory that is asserted, and the conceptual framework used to assert it. On the contrary, I have argued throughout that the distinction between conceptual and empirical propositions is one of function. It has never relied upon the criteria we use for adopting or revising different principles. Empirical facts may refute certain empirical principles we hold, but conceptual frameworks are not the kind of thing to be refuted. That does not prevent them from being rendered inapt or useless by empirical facts, though this process is mediated by whatever interests we have. Whether a lexicon is any good will depend on how it fits the purposes for which it is designed. A set of scientific terms are naturally evaluated according to their power to articulate adequate theories. Hence, the criteria for adopting a particular theory or conceptual framework are often indistinguishable.[311]

In this way, the relevant distinction should not be sought in the criteria used for adoption and revision, but rather in the functions different principles have. As Friedman writes:

> What characterizes the distinguished elements of our theories is rather their special constitutive function: the function of making the precise mathematical formulation and empirical application of the theories in question first possible. (Friedman 2000, 377)

That some principles constitute what can be said with a given set of terms, while others say things with those terms, is how the distinction must be maintained. Constitutive principles determine meaning and give language its expressive power. If we stick with Einstein's revolution in physics as our example, 'the whole notion of a variably curved geometry itself only makes sense in the context of the revolutionary new theory of manifolds recently created by Riemann.'[312] One could hardly discuss the idea of a variably curved spacetime in centuries previous, let alone use it to describe empirical phenomena.[313]

4. Semantically variant theories

My response thus far suggests conceptual change is an aid to scientific progress, not an obstacle. But it is worth returning to the original problem to see if it can be answered more directly. One

[311] Cf. Carnap (1950, 23-4).
[312] Friedman (2000, 376).
[313] Cf. Kuhn (1993, 331-2).

way of expressing Putnam's objection is the following. The claims of competing theories cannot conflict with one another if it turns out they use different concepts, because they will be talking about different things. Indeed, how could we be said to learn new things about something when that something is not the same across different theories?

I have already provided one answer to this question. Conceptual change need not be radical, so the notion that we are talking about completely different things is a misapprehension. But even radically different concepts can describe the same phenomenon: 'the box on your desk' versus 'your computer'. Conceptual difference doesn't require different extensions, and sameness of reference allows content to be compared.[314] Even partially shared extensions still allow theories to make claims over (some of) the same objects. Theories using such concepts may be incompatible based on what they say regarding the shared objects referred to (and thus one may be found to better accord with inductive evidence).[315] Indeed, one feature of non-radical conceptual change is, presumably, that there is significant overlap between the relevant concepts' extensions.[316] In some cases, the extensions may be exactly alike (or alike in all known cases). Israel Scheffler makes essentially the same point, drawing on Frege's distinction between sense and reference.[317] Of course, pure referentialist theories of meaning have no truck with intensions, but I have already given reasons for rejecting those accounts. At a bare minimum, reference-fixing requires some conceptual content.[318] This could include a description of the object's causal role, a categorisation of the object, or the outlining of some kind-constituting properties.

Secondly, even where there is no referential overlap between the technical terms of different theories, it doesn't follow that the two cannot be compared. There may be a shared reference point outside of those technical vocabularies. Where two theories are said to be semantically incommensurable, they are not to be understood as distinct languages. Theories will arise from a wider, natural language that incorporates both theoretical vocabularies. Semantically variant theories are merely theories whose conceptual components are not translatable into each other.[319] There is no one-to-one comparison: they carve logical space differently. The sense that

---

[314] Cf. Sankey (1994, 38-41; 2011).

[315] Devitt (1991, 170) describes several 'quasi-logical relations' that can exist between different concepts with referential (or extensional) overlap.

[316] Note that it needn't be a one-to-one replacement of concepts with similar extensions. Consider a case where what used to be described as an F is now described as a G that is H and I, which shows how new concepts to describe the same phenomena can be entirely different.

[317] Scheffler (1967).

[318] See chapter three.

[319] Davidson (1973), of course, argues against the notion of distinct conceptual schemes. See Hacker (1996b), Glock (2007) and Sankey (1997) for convincing critiques. Neither the claim that distinct conceptual schemes must be untranslatable, nor that we couldn't recognise a practice as an untranslatable language, seem to survive scrutiny. The Russian colour-scheme or the pre-industrial way of measuring time in Japan, for example, certainly seem different

these theories are 'incommensurable', as Sankey notes, 'is due to semantic differences in the terminology of theories: the terminology employed by a theory cannot be translated into the terminology of a theory with which it is incommensurable.'[320]

This helps because, although it may be true that the entities different theories refer to vary, it may nevertheless be the case that a different kind of shared reference exists – one that can be appreciated in the wider language. For instance, two theories may be trying to explain the same phenomenon, or it may be one implication of the theories that the same phenomenon is to be explained in different ways. Combustion, say, or the movement of celestial bodies. Alternatively, theorists might be giving an account of a substance that is said to be isolated by a particular experiment (and both theorists may accept this).[321] They may explain what they think the isolated substance is and develop a conceptual framework with different components to explain how the substance becomes isolated by that experiment. Someone else might develop a separate theory, with distinct conceptual components, outlining a different causal mechanism. These theories need not share concepts for them to be compared. After all, if I ask you to guess what is in my hand, your answer can be compared with someone else's regardless of whether you've used the same concept (and hence guessed the same thing). Shared references, which might act as an anchor for comparison, needn't be a feature of the narrow explanations, theories, or technical vocabularies themselves. Even in cases where there is no extensional overlap with respect to those technical terms, they may nevertheless be used in different theories that would, if true, explain the *same* phenomena. In such cases, it might be the wider language that allows us to appreciate these broader (in some cases, macro-) phenomena that are the same across different theories. Theories can thus be seen as competing even if the technical terms those theories employ are designed to refer to different entities (which may or may not exist).

This discussion gives sense to the idea of *seeing the same world differently* or *thinking about the world in different terms*. Both phrases are associated with moments of significant conceptual change. And both imply that there is a globally shared reference point. The thought is that our empirical

---

from our equivalent schemes. Yet they can be translated. And we can recognise a practice as exhibiting characteristics typical of language without having any understanding of it (because, perhaps, speakers refer to samples we don't have access to, or they have perceptual capacities we lack). This may be in tension with §207 of *Philosophical Investigations*. Wittgenstein claims of an activity resembling language that we may judge it as lacking the regularity necessary to count as 'language'. But is there any reason to think a language must be learnable by us, or to think we are equipped to recognise regularity in all other languages? The story does, after all, provide plausible evidence of communication, namely the behavioural changes that occur when members of the tribe are gagged. More generally, there is evidence we could have that wouldn't presuppose our ability to recognise regularity in the language itself. We might, for example, observe the teaching of a language, where it is clear from the teacher's behaviour when they are issuing corrections or praising their pupils.

[320] Sankey (1997, 86).

[321] Cf. Burian (1985).

concepts are tied to the environment we find ourselves in and created to describe it. Of course, when looking to compare theories, a shared reference point of *the world* may be of little use, but there may be narrower shared reference points connecting competing theories. Narrower reference points could include phenomena or processes that might be referred to using words that are not technical terms particular to a theory's vocabulary.

5. Conclusion

There is, of course, much more to be said on this, but it goes beyond the scope of this chapter and thesis. The aim of this chapter has been to show that it is not obvious that scientific discourse is problematic for my account and that some arguments designed to establish that conclusion are unconvincing.

This chapter has also raised a wider question: namely, what is the relationship between the world known through experience, and that is the subject of empirical propositions, and conceptual propositions. The arguments above suggest the relationship is close, since they imply that our conceptual resources are shaped by the world and our knowledge of, and beliefs about, it. The close interaction between the world and our conceptual resources is one source of the Quinean objection we saw Maddy outline. It has also provoked others. Part III of this thesis considers more generally the relationship between reality and our conceptual resources.

(8) Identifying Conceptual Norms and Criteria

Quinean doubts about meaning and the distinction between conceptual and empirical questions often arise from epistemological concerns. In effect, the question is: can we ever – in practice – successfully distinguish between conceptual and empirical propositions? The answer previously given was that there must be such a distinction in language and that it is manifest in the difference between surprise and incomprehension. I argued that something must hold fast. Conceptual connections are the pivot upon which empirical discourse turns. For us to say anything at all, and to make empirical judgements, it must be determined what counts as a phenomenon of the kind about which we (attempt to) speak. This much must be settled.

This chapter is intended to suggest some methods for making explicit our linguistic norms. I suggest several ways for us to determine that a given feature acts as a constitutive criterion. While I agree that indeterminacy is widespread, I nevertheless think there are definitive examples beyond acts of pure stipulation. We can identify conceptual propositions at least some of the time.

1. The epistemological conception of analyticity

Let us start by stating some clear examples of the phenomenon we want methods to identify. 'Blue is a colour', '"better" is a word' and 'every father has a child' are all conceptual propositions. One does not know what 'blue' or 'father' means without accepting the relevant propositions. Likewise, one does not understand the sequence of letters 'better' as used by us without recognising it as a word. These cases seem trivial but given we are free to acknowledge them they do at least push against the kind of scepticism that apparently haunts us. For example, having a child cannot be understood as an invariably common symptom of fatherhood. It is not just that, as it happens, there are no fathers without children but that there *cannot be* fathers without children. We would not be surprised by someone who claimed to have discovered such a 'father' but bemused by what they had said. This is not to say that should someone claim that a father need not have a child we would have to conclude they were not making sense, or they must not have the same concept 'father' as us. We might find that what they meant was that one does not need a living child to be a father, or that one is still a father when one's child becomes an adult. In these cases, they are not attempting to deny a conceptual proposition, they simply have something different in mind when using the statement that we would otherwise use to express a conceptual proposition.

Moreover, even in circumstances where one correctly interprets the conceptual proposition, one's refusal to accept it does not entail that one has a different concept altogether, or that others will fail to understand one because of one's refusal. For 'understanding' comes in degrees, and we often use words without knowing the precise conditions under which they are correctly employed. One can happily use the word 'whale', for example, without knowing that whales are a type of mammal, as some young children plausibly do. Here we can accept that they are trying to use our concept, only their understanding is less advanced than our own. Concept possession comes in degrees and one can be said to possess a concept (albeit perhaps not fully) without a total mastery of its application conditions. One can use a tool without being a master of it.[322] But it seems a bit of a stretch to say that they could share our concept 'whale' without knowing that it is an animal or living thing. If they think stones can be whales, then we might think they have not yet grasped the concept sufficiently well for us to say they possess, or understand, it.

There can be differing degrees of self-consciousness with respect to one's own mastery. The child may not know that there is a more detailed definition of the concept they use, whereas many mature language-users do make use of words that have technical definitions that they know they are not in possession of. They are, therefore, willing to acquiescence in many such cases to those with the technical expertise should a controversy arise. Now, if a mature language-user has no interest in acquiescing to those with technical knowledge, we may eventually accept they use a different concept from the expert or, at any rate, wish to – whether they are consistent in applying their concept according to their lay standard would be something to investigate.

This discussion hints at a criterion for conceptual propositions, one that can be employed as a method for determining whether a given statement qualifies. Interestingly, it is rather close to a definition of analyticity that Quine himself considers, namely: 'a sentence is analytic if everybody learns that it is true by learning its words'.[323] This is somewhat close to the epistemological conception of analyticity that has been proposed throughout the twentieth century, which we met (briefly) in chapter six.[324] Now, as Glock rightly suggests, 'Quine's definition includes a wanton genetic element that is easily jettisoned' and has been by many

[322] The distinction between possession and mastery lies at the heart of recent exchanges between Boghossian and Williamson (2020, e.g. 245-6) on the nature of the *a priori*.
[323] Quine (1974, 79; see 1986, 93–5).
[324] For instance, in Grice and Strawson (1956) and Boghossian (1996).

proponents of the epistemological conception.[325] Thus, we arrive at Glock's favoured formulation:

> … if a speaker x sincerely denies or rejects s [a sentence expressing an analytic proposition], this shows either that x fails to understand s, or that x is deliberately employing s in a novel sense. (Glock 2003a, 83)

In light of the discussion regarding degrees of understanding, we might like to qualify Glock's proposal. For it may not be tantamount to a failure to understand so much as a failure to understand *fully* or *completely*. In other words, it is possible that someone has a limited grasp of the relevant proposition, and the concepts contained within it, without their having the kind of understanding that would lead them to accept a statement of the linguistic norm. In those cases where the negation of a proposition can only be explained by reference to a misunderstanding (of some degree) or a different use of familiar words, it will be conceptual. A mark of a conceptual proposition is that one or more of the words must change in meaning for its apparent statement to be denied. This remains consistent with what I have said regarding degrees of understanding, because we are here taking the child's limited understanding of 'whale' to be a separate meaning for the term ('large animals that live in the sea', perhaps). When the child denies that whales are mammals, they employ the term 'whale' according to a simplified (and, therefore, different) standard. While the child's understanding is really an impoverished version of our own (an attempt to grasp it), it is also a possible explanation for that term. We are, therefore, able to acknowledge both truths: on the one hand, children are trying to learn our language and their conceptual repertoire is a simplified version of our own; on the other, their limited grasp of our concepts, and the standards they employ, could be used to define a different concept.

2. Possible cases

The epistemological conception of analyticity suggests one possible method for identifying norms, though it could not be relied upon in all cases. Where errors are less egregious, and more subtle, reactions may vary. In cases where norms are imprecise, the distinction between

[325] Glock (2003a, 82). He points out that it is the genetic component of this conception of analyticity that is one of Quine's complaints (1976, 119–21; 1986, 95, 138, 206). Thus, one of Quine's principal qualms with his attempt at a definition of analyticity is easily disposed of.

incomprehension and surprise may blur. There is, however, a further method we can rely upon at this stage. How the different statements are, or can be, defended is another possible route.

> Suppose it were said that 'An unmarried male is a bachelor' is true by definition. Many people would accept this as correct, not noticing that an unmarried male of five years of age would not be called a bachelor. How would it be shown that they are mistaken? Not by empirical evidence (a survey of bachelors, etc.), but by reference to language – in this case, the use of 'bachelor': we would point out that this word is not applied to children. By contrast the statement 'Bachelors are more wealthy than other men' would have to be investigated by empirical methods. (Hanfling 2000, 217)

When defending a conceptual proposition, one considers how we do (and would) use our language in a range of different cases. We look to possible cases to show when it is appropriate to apply the relevant words. In the empirical case, we take for granted the objects we are referring to and investigate whether those objects are as the proposition describes. Likewise, if someone makes a statement and we do not know whether they mean it as a definition or description, asking them to explain how they know what they have said may produce results.

From this discussion, another related method for establishing conceptual propositions suggests itself, one that neither relies on what we say when asked to explain the meaning of a given word, nor our justification for using that word (both of which are themselves legitimate methods of inquiry). We could begin to establish what the relevant norms are by asking competent speakers to respond to an array of different cases. The task would be to identify cases in which the relevant concept applies. This is a way of establishing where those with an implicit understanding of the language draw the line between correct and incorrect uses.[326] Even without asking them why they are excluding a particular set of cases, these reactions can inform us. From these reactions we could begin to build hypotheses as to what different words mean, and we could design further cases to test our hypotheses.[327] This strategy is endorsed by Carnap, in response to Quine's scepticism about intensions.[328]

3. Explanation and justification

[326] See the final section of this chapter for a discussion that touches on some of the limitations of this method. It should be noted that I suggest this as a way of *beginning* to establish norms. Inside our own language, it is a peculiar method if isolated from explanations and justifications.
[327] Cf. Thomasson (2021, 2091-2).
[328] Carnap (1955).

As mentioned above, another method involves simply asking for explanations as to what is meant by something. Even in cases where I identify what I mean with an empirical description, I will tend to clarify the category of that something. For example, when asked for the meaning of 'cheetah' I might reply 'it refers to the large spotted cats found in Africa' and for 'lettuce' I might say 'it is the name given to leafy plants like that grown in our garden which we add to salads'.[329]

A related method concerns the justification for a manner of speaking. If pressed for a justification for the use of some word (rather than the evidence for a claim), I will justify my use by rendering explicit some of the norms that determine the relevant concepts' employment.[330] This can be done directly by appealing to the norms themselves, or indirectly via facts about what competent speakers do. One might reasonably wonder how we can tell whether someone is asking about the evidence, and to be clear it may not be obvious from the start. But suppose I clarify by stating my evidence, my interlocutor may press me further by saying they accept all of that and are rather interested in why I think that the thing it points to is the phenomenon I claim it to be. Suppose they do not doubt the features that are present, only the classification of them. At this point, one's appeal will be to what phenomena of that kind are (their identity). The way to adjudicate over these matters is to consider what counts as that kind of phenomenon, thereby appealing to conceptual propositions. At the bottom of our explanatory chain are rules for the use of words: 'this is what we call. . .', 'this *just is* a. . .', 'this is what we mean by. . .'.[331]

What enables explanation to end with constitutive norms? Consider 'all proper names are spelt with a capital letter'. To spell proper names correctly requires that one does as the norm says. But no further explanation (of the norm itself) is required. Once one has fulfilled the relevant conditions, then one is engaged in the practice constituted by that norm (because engaging in that practice *just is* following what the norm requires).

4. Necessity

We have now returned to familiar ground. What is distinctive about conceptual propositions is that they can be expressed in statements that incorporate a logical *must*. This marks a constitutive relation, typical of semantic norms that set out conceptual criteria.

[329] It is the partly empirical nature of these definitions that makes Waismann (1950, 30) resist the connection between dictionary definitions and analyticity.
[330] Cf. Waismann (1965, 141). See also Sultanescu (2022) on rule-following and rationality.
[331] Cf. Friedman (2000, 383) in a different context: 'our knowledge *has* foundations in the present sense: subject-defining or constitutive paradigms'.

We can of course apply words according to symptoms. I can use symptoms to identify, for example, certain types of animals or flowers as I stroll through countryside. This is a perfectly rational procedure. Symptoms are good, empirical evidence, but they are not definitive. Even defeasible criteria provide what we might call 'logically good evidence', which is intended to acknowledge that the relation between the phenomenon we have identified and one of its features is not merely an empirical association (even if – on its own – it is not definitive). This is the difference between meeting one of the standards and having qualities which indicate that you might meet them. Consider the following example. While having a fever, cough and fatigue might suggest you have been infected by a flu virus, none are required to be infected. Indeed, none are criteria. The relevant criterion is whether a virus of the correct type has invaded, and the extent to which it has multiplied inside, your body. Whereas 'having the flu' is generally understood as having the illness caused by the relevant virus and so does require symptomatic disease.

To be clear, the 'must' can flow in both directions. Some phenomena undoubtedly do have criteria that must be present for something to count as such. But many plausibly do not. In those cases, the 'must' is more easily observed in the opposite direction: *if this set of features is realised, then it is that phenomenon*. Here there is no room for doubt. Where this set of features is realised, even if it does not have to be this particular set, the phenomenon in question is necessarily of that kind. This is a matter of linguistic legislation. In theory, an exhaustive disjunctive list of such sets could answer the question of what a phenomenon of that kind must have. But I have already raised doubts regarding the plausibility of distilling such exhaustive lists with any kind of precision in our natural language. Moreover, family-resemblance concepts may resist such treatment.

Answers to questions that employ 'must' can, therefore, point us in the right direction. 'What must a good friend be like?' or 'what is required of a good friend?' ought to provoke responses that provide conceptual guides. They might say 'a good friend must be. . .', 'a good friend must either be. . . or have. . .', or 'a good friend must combine. . . and . . . or be. . .'. These formulations are intended to show how such questions can be answered with both criteria that are individually necessary and those that are not. As should be clear from previous chapters, one would not expect a set of precise conditions, though perhaps there are some people who define 'good friends' according to strict standards. For most of us, we would be able to provide an inexhaustive list of the kinds of qualities required, and perhaps some examples of the things a good friend would do. But others might have their own idiosyncratic ways of determining such matters. Like ours, but not the same.

5. Choice

Of course, we could adopt their standard, and this fact suggests another mark of conceptual propositions. We can *choose* to accept or reject them. It is a matter of fact as to whether we do or not, but upon encountering conceptual propositions we have a choice in a way that we don't with matters of fact. To be sure, one can dispute a statement of fact by rejecting definitions upon which it relies, but there one is not denying the fact but rejecting how it has been described. Keith Donnellan provides the following example.[332] Suppose we discover that some of the creatures we have thought of as whales turn out to lack the features necessary to be considered a mammal. One person might claim: 'see, not all whales are mammals after all'. Whereas I would be inclined to say: we have discovered that these creatures are not really whales. Here we are agreed on what has been discovered, namely there are creatures that have been called 'whales' in the past that turn out not to be mammals. The difference concerns whether they can continue to be classed as whales. The other person says 'yes' and I say 'no'. But both of us will surely recognise that we may choose to go either way. I can, for instance, momentarily adopt the other person's understanding of 'whale' and accept that what they say is true. In fact, this is an infelicitous way of putting it. I need not adopt their concept at all. I can simply accept that, on their understanding, what they say is true. We might have reasons for preferring our concept 'whale': we may be sure this is what we previously meant by 'whale', or we might think this is the better way of categorising marine animals. This is incidental to the point being made, however. Whatever can be said for or against the different concepts 'whale' we are proposing, it is fundamentally a *decision* to stick with one over the other. It is possible to choose otherwise, in a way that there is no such choice regarding the discovery. This element of choice marks conceptual propositions from their empirical counterparts.[333]

Donnellan contrasts this case to one that involves the proposition 'all cats drink milk'. Suppose I think all cats drink milk while the other person disagrees. Let us further assume this is understood as an empirical claim about all cats. Upon finding a cat that does not drink milk,

---

[332] Donnellan (1962, 656-7).

[333] Sidelle (2007, 92) provides a similar example: 'Melville makes it quite clear that he understands Linnaeus's motivations for not considering whales fish, and he allows that whales are warm-blooded, breathe air, bear and suckle live young, etc., in common with other mammals and not with other fish. He simply disagrees that this suffices to keep them from being fish. Similarly, there need be no disagreement about the intrinsic features of bats and (other?) birds, their similarities and differences both with each other, and with other natural creatures. But aside from these matters, upon what can the truth of the issue in dispute depend? The most – I propose – is that they can depend upon the established use of the words "fish" and "bird" – which is not a matter of the nature of fish or birds, or any other biological matter, but a matter, at best, of semantic history.'

neither of us will be inclined to rule that this creature is – based on its drinking habits – not a cat. Yet I might still come up with further explanations to defend the claim that all cats drink milk. For example, I might suggest the cat drinks milk secretly when we are not watching. But this way of avoiding the other's conclusion is not analogous to the conceptual case where we refuse to employ the other's concept. As Donnellan writes:

> It may be true that I can hold on to the statement that all cats drink milk by explaining away every apparently falsifying experience. Eventually I may have to describe some as hallucinations or illusions. But the two of us are not faced with a choice as to whose way of describing the situation is to be used. (Donnellan 1962, 657)

The moves we make, and the justifications we call upon, differ depending on the kind of disagreement we are involved in. In the empirical case, the content of our disagreement is not what counts as a cat drinking milk, but whether every cat does, in fact, drink milk. If one can be said to choose to believe an empirical proposition despite all contrary evidence, it requires something quite different from the conceptual case. Contrary evidence must be explained away. We might choose to believe (or keep believing) something, but we accept those beliefs can – in principle – be falsified. An empirical statement requires that we can specify under what circumstances it would be false. This is the sense in which we cannot 'choose' to go either way. One of us will be correct. If this can be proven, the other will have no choice but to accept it. If 'all cats drink milk' is a conceptual statement, there are no circumstances in which it can be false. Where we have semantic differences, we can acknowledge that both of us are correct (given our different choices).

To say they are choices does not, however, mean that abandoning certain norms would not be of great consequence. The choices we make are not costless. Not all choices are 'equally good'. Some norms might be central to our way of life, the way we think and act. They may not be easily discarded. We might be unrecognisable without them. Some principles of logic, for example, are intimately connected with what we call 'thinking' and 'reasoning', and we may not be able to think or reason without them. Moreover, to say that we have some choices now does not mean we chose the norms in the first place. Nor does it mean we have a great range of choices in every case. The relation between norms and facts (and other matters that might influence our 'choices') is the subject of the final part of this thesis.

6. 'The paradox of analysis'

We have been exploring ways to establish conceptual norms and criteria. Before ending this chapter, we should consider two objections that relate to that discussion. The first concerns the supposed 'paradox of analysis'. It might be said that the methods my account relies upon presuppose what they seek to find out, since the answers are, in effect, articulations of the norms (or the logical consequences of norms) we already understand and follow. That is not a problem in and of itself. In fact, it might make one confident that the account has a plausible epistemology. However, some argue that this confidence comes at the cost of making the exercise it describes trivial. If it were worthwhile, then the answers would be informative. But how can they be informative when we already understand the norms that we are seeking to uncover? Here we encounter a version of the supposed 'paradox of analysis'.[334] How can we learn something from a process that presupposes knowledge of that which we are supposed to learn?

To resolve this paradox, we need to recognise different types of understanding. We might consider the process a kind of distillation, for it is a matter of separating an explicit rendering of linguistic norms from our actual, everyday use (which may involve explicit mentioning of certain norms but does not have to).[335] Norms may only be latent in use, but our aim is to make them patent. Much of our language-use is instinctive and we do not as a matter of course consult rules. We speak naturally and it can often seem automatic. Here we may employ the language of Mark Balaguer and Terry Horgan and say that one can have an implicit and explicit knowledge of the meaning of a term.[336] An implicit knowledge, or understanding, gives one what we might call 'application competence', namely the ability to apply the word in ordinary circumstances without making mistakes. An explicit knowledge entails something different: the ability to provide a definition for a word (if one can be given), or to explain its various conceptual connections. This latter ability enables us to provide an explicit rendering of the relevant norms and their consequences.[337]

To engage in our method of analysis requires implicit understanding, but it aims to deliver an explicit one. This is by no means intended to imply that we could not interpret a language that we had not been trained in. In fact, the method in such cases is not altogether different: we study the circumstances in which competent speakers apply their words, their

[334] One might see this as a variation of Meno's paradox, where Meno wonders whether it is so much as possible to conduct an inquiry into the meaning of 'virtue' (*Meno* 80d-e). See Hanfling (2000, 58).
[335] Glock (2012b, 65-7) is relevant here: someone might recognise a rule as the one they follow without being able to articulate it first themselves.
[336] Balaguer and Horgan (2016).
[337] Cf. PI, §109.

reactions to various sayings, and approximate from their behaviour what they mean. In other words, we look at how people who have an implicit understanding of the language act and try to distil the norms by which they apply their words. The reason for not paying too much attention here to languages we do not already understand is that the charge of triviality is not relevant since we would be learning a new language.

Anyone familiar with sport ought to recognise this overall response as having force. It is common for someone to play a sport well without being able to explain how they do so. There is a difference between showing someone how to do something and being able to explain how to do it. For instance, one may themselves have a beautiful cover drive but be unable to explain the fundamentals of one to a young cricketer. That is even the case where they might have first learnt to play via explicit coaching. Some sportspeople will have no trouble whatsoever explaining the basis of their technique, while others play sport instinctively: they have a *feel* for the game. Such players might be able to identify when something has gone awry with their technique ('it doesn't feel right') but even after watching a video may not recognise where the problem lies. This distinction is not binary, and many of the best players will have some combination of feel and explicit understanding. But the point here is to suggest that the distinction proposed in the case of language is also present elsewhere.

7. A methodological question

The second objection concerns a controversy that has long been associated with philosophy and especially those conceptions of the subject that give language a central role. Sometimes these are posed as challenges to the role of 'intuitions'. Why should we trust the intuitions of philosophers? What gives them privileged access? While I have afforded no special role to 'intuitions' and they can therefore be left to one side, linguistic philosophy can be probed on a related basis. What gives philosophers the right to determine how *we* speak? Are we not obliged to ask the community to whom 'we' refers? The question is whether it is so much as possible to determine linguistic norms from reflection, or whether it must instead be done by empirical research. Less radically, we might ask whether empirical investigations might be more reliable than armchair study.[338]

[338] As Hanfling (2000, 4) points out, a contributing factor to the scepticism regarding this philosophical method was that 'rather obvious mistakes about "what we would say"' were made by its leading protagonists. Hanfling (*ibid*, 56) discusses this in relation to Ryle's (1949) and Austin's (1957) discussions of 'voluntary' and 'involuntary', something Cavell ([1958] 2002) also tackles without obviously getting it right (cf. Sandis 2021). Mates (1958, 165) even suggests

This challenge is historical, though it has its contemporary counterparts.[339] One of its early iterations came about when Cavell was challenged by his colleagues to defend the methods of linguistic or, more specifically, ordinary language philosophy.[340] His opponent was Benson Mates who, along with advocates of 'empirical semantics', were sceptical of Austin's method for investigating ordinary language.[341] While Cavell was a great admirer of Austin, they rejected Austin's armchair philosophy and thought linguistic surveys of competent speakers ought to replace it.

7.1 Answering the challenge

It is important to recognise that procedures for identifying norms are no less fallible than empirical investigations. Indeed, in some cases they may be akin to those investigations. However, when we are competent speakers ourselves, this description is not quite right. We no more need to consult evidence to recognise an incorrect linguistic move than we need to observe how rugby is played to know that one may not pass the ball forward. Similarly, the rules of chess entail that a pawn must have taken an opposition piece for two pawns of the same colour to be standing in the same column. Hence, I can have *a priori* knowledge of this necessity without having played, or watched, a single game of chess. Part of what it is to be a competent speaker, to understand our language, is to know of these connections.[342]

Fluent speakers of a language are such because they understand the norms by which their language is spoken. These they 'learn from parents, teachers, rule-books, dictionaries, literature etc. as well as by observing general linguistic behaviour, appealing to precedents to resolve cases

---

the mere existence of disagreement between Oxford Professors is evidence enough that this method is doomed to fail.

[339] Contemporary counterparts typically fall under the heading 'experimental philosophy'. See, e.g., Appiah (2008), Knobe (2003), Knobe and Nichols (2008), and Weinberg, Nichols, and Stich (2001). But note Knobe's (2007) response to Kauppinen (2007) where he suggests experimental philosophy (at least as he understands it) is not designed to answer semantic (or conceptual) questions at all. It is rather a return to a more traditional (we might say, metaphysical) conception of philosophy. This is a more fundamental disagreement, and it is not one I tackle here. Appiah (2008, VI), on the other hand, does seem to make some of the mistakes I will explore below. For instance, he seems to conflate use with usage, insofar as he thinks empirical studies, that tell us the percentage of people who respond in a certain way to a given case, are better placed to answer conceptual questions than the methods linguistic philosophers typically employ. See Sandis (2021, 308-9) for criticism.

[340] See Hansen's (2017) introduction for more of this history. He takes much of it from Cavell (2010).

[341] Mates (1958). Naess (1957) is a classic example of empirical semantics of the relevant kind (see also his 1938). Chapman (2011) is an introduction to Naess's philosophy of language. See Hansen (2017) and Sandis (2021) for slightly different accounts of why Mates' response to Cavell might be mistaken. Fodor and Katz (1963) also pose objections to Cavell. See Bates and Cohen (1972) for a rebuttal.

[342] This defence (and the analogy with games) is a common response to this problem. See Sandis (2021), Searle (1969), and Vendler (1967).

of conflict.'[343] To establish what those norms are, then, requires them to reflect on the ability they have. It is not mysterious that they can establish linguistic norms without recourse to empirical research. Indeed, it is not obvious that they need consult evidence at all. They reflect on what it is they do and try to render that understanding explicit (to varying degrees of success). This answer relies on the distinction between different modes of understanding we saw earlier. An implicit understanding is the ability to apply a word correctly, while an explicit one implies that one can articulate the relevant norms. Engagement in a process of conceptual elucidation requires implicit understanding but aims to deliver an explicit one.

Our primary concern is whether linguistic philosophers' means of arriving at this explicit understanding is appropriate, or whether it ought to be replaced by a more robustly empirical method. What would an empirical investigation of the relevant kind look like? In general terms, it would be a matter of asking fluent speakers how they use, or understand, various words. It would, therefore, rely on the very method some seem to think it could replace: fluent speakers reflecting on how they use their words. And this points to a further reason, on top of the fact that competent speakers are initiated into a norm governed practice when they learn a language (and hence have access to those norms), why fluent speakers need not rely on evidence:

> [Native] speakers do not, in *general*, need evidence for what is said in the language; they are the source of such evidence. (Cavell 1958 [2002], 4)

Sandis' reflections on this passage are apt:

> When an individual native speaker fails to notice difference or similarity in linguistic usage, the evidence that tells against her analysis is proffered by other native speakers. The important point is that the evidence does not lie outside of the community of native speakers in such a way that would allow for a foreign ethnographer to amass it and thus prove that all native speakers were mistaken. (Sandis 2021, 302)

This is not to rule out the usefulness of empirical research altogether. Just because one does not typically need evidence does not mean it cannot be useful. Such procedures could, for example, remind us of a use we have ignored, as could our reading of a newspaper, novel, or employment contract. But empirical investigations can have significant shortcomings. For one, overreliance on them puts at risk the distinction between correct and incorrect uses, by conflating use with usage. Let the former be a normative notion, while the latter concerns custom, habit, or

[343] Sandis (2010, 186).

disposition.[344] Use marks what it is correct to say, while usage represents what we do, in fact, say. Usage is certainly a helpful guide to correct use, but that does not mean one exhausts the other.

Antti Kauppinen's gloss on phrases, often used by linguistic philosophers, like 'we would ordinarily say', 'we are inclined to say' and 'we usually say' is helpful here. They are best understood as elliptical: it is what we would say provided we i) are competent users (or better, masters) of the relevant concepts, ii) have considered the case in sufficiently ideal conditions, and iii) are influenced only by semantic considerations (rather than, say, pragmatic ones).[345]

Competent speakers do, after all, get things wrong. In conversation, we often accept corrections from others – acknowledging that we have misapplied a word according to the norms we otherwise follow. The distinction between correct use and actual usage does not imply that an old-school grammarian is fit to dictate to one and all when words have been misapplied. We employ this distinction even within the context of our own idiosyncrasies. This distinction is what makes it that words mean one thing rather than another. There are standards that explain what those idiosyncratic uses mean. What one does can be measured against those standards, regardless of how widespread they are. In criticising the self-description of someone's idiosyncratic use of a word, we do not mean to say it doesn't accord with established usage, but rather that it fails to 'accord with the practice of the person giving the description.'[346] The person giving the description has got something wrong about themself. There is no philosophical problem with using words differently (so long as we do so consistently). In other circumstances, it may be entirely appropriate to criticise one's use of a word on the basis that it is non-standard. These could include conversations where others use the word standardly, and so by treating the word in a non-standard way one misconstrues the point being made, or if one is trying (but failing) to speak as others do.

Nat Hansen criticises the part of this response that rests on the distinction between use and usage, by pointing out that philosophers have no special access to the relevant facts about correct use.[347] Why should their reflections have priority over anyone else's? But the claim Hansen contests need not be made. Each of us may challenge rule-descriptions as we see fit. If the philosopher has an advantage, it is not one of privileged epistemic access. Rather, philosophical skill resides in elucidatory and cartographic capabilities.[348] The ability to recognise

[344] This is a useful distinction, though English does not always respect its boundaries. Sometimes 'usage' is synonymous with 'correct use'.
[345] Kauppinen (2007, 101).
[346] RPP I, §548.
[347] Hansen (2017, 794-5).
[348] Sandis (2021, nn. 76).

and make perspicuous conceptual and logical connections. The goal often being to use these perspicuous representations to resolve philosophical problems.[349] To achieve this, it is irrelevant whether the whole community uses the relevant concepts in the same way. All that matters is that we successfully describe the concepts used to formulate the philosophical problem under consideration.

Furthermore, what special advantage is empirical research supposed to have? It is entirely plausible that a straightforward poll that asked people to react to different cases by saying whether they fit a certain description would be a reliable route to the empirical facts about how those people speak. But could the same be said for measuring correct use? In many cases: yes. Fluent speakers are not, in general, wrong about the language they use (otherwise they would not be counted as such).[350] However, this answer must come with the following qualifications.

Firstly, we often *do* require linguistic corrections (from ourselves and each other). Our task is not always a straightforward one and the cases we are typically interested in when doing philosophy are precisely those that pose significant challenge. A reliable method in that case requires more than the mere recording of people's reactions: it requires challenges to be made, explanations to be given and further cases to be reflected upon. This is not impossible for empirical research to incorporate, of course. But the relevant criterion of correctness would still be the standards of our established practice, in the same way that the answer to '103 × 5' is dependent on the rules for multiplication, not the answers of a surveyed majority (even if that majority is likely to arrive at the correct answer).[351] For that reason, individual contributions to a survey would have to be judged based on what was said in response to the challenges raised, for participants might have spurious reasons for their answers. A simple majority verdict would not be conclusive. It would merely record what the majority thought the right answer was. Perhaps a different interlocutor might have been able to convince them otherwise. Of course, we may also

[349] Cf. PI §122.
[350] Cf. PPF §348 where Wittgenstein calls into question the sense of everyone following a rule incorrectly. Stroll (1994, 112) also discusses this in a different context. See chapter four.
[351] This analogy shouldn't be stretched too far. Language is not a calculus, as Wittgenstein very well understood (cf. BB 25 and PI §81). The rules in mathematics are generally clear, precise, and formally laid out. Whereas linguistic norms (with exceptions) are far more elusive, often in flux they can even have an almost evanescent quality to them. But there is usually enough stability to be able to give some characterisation of the criteria we use, or the norms we follow (which may be no more than pointing to the relevant family resemblance). Without this stability, communication between us would be compromised. We would not know what each other were saying. It is true that our norm-, or rule-, descriptions may often be vague (or open-ended), lacking necessary and sufficient conditions, but that is only a problem if the relevant norms are really more precise than our characterisations suggest. See Mulhall (2001) for a discussion regarding the differences between calculus and language.

uncover a number of untypical, idiosyncratic uses of familiar terms which do not stand in need of correction.

Notice how this augmented research method increasingly comes to look like something which is recognisable as philosophical dialogue.[352] Seen in this light, the positive claim for empirical research becomes something like 'we would do better philosophy were more people doing philosophy'. Even if this were true, it sheds little light on philosophical method. Moreover, it is perfectly reasonable to think that philosophical training might (and probably *should*) make its participants well-equipped to engage in this kind of inquiry: to ask the right questions, identify challenging cases, and proffer clear explanations.

7.2 An unanswered question

Whatever the merits of this response, however, it does risk obscuring a matter still to be settled. For there remains an empirical question to be answered even if we accept the distinction between use and usage, and hence accept some of the cautionary notes about empirical methods devised to establish linguistic norms. Moreover, it is one Hansen identifies in his article. When discussing the analogy with games, he points out that we are still required to make a separate claim about the relevant rules being *shared*, and that is not something to be known *a priori*. The exact extent to which they are shared (and by whom) is something to be investigated empirically. It might be right that to understand the norms being followed requires the kind of dialogue mentioned above. That is, it might demand a robust form of self-reflection with respect to how one speaks. But while that might be cast as an *a priori* exercise, the point is that an empirical question can be asked about those exercises: namely, 'which people follow what set of rules?' And this question seems difficult to divorce from the linguistic philosophers' conclusions about how *we* speak.

To be clear, the distinction between correct use and usage is not threatened by the possibility of others following different rules, as the analogy with games should make plain. If rules are not shared, then we are not playing the same game. But someone's playing a different game hardly contests the fact that the rules of *this* game are as we say. No one is obliged to play chess by the rules we do. Some might play a variant of it. But that wouldn't be our game of chess. If someone were to move a pawn sideways in our game, for example, they would be told they cannot do that. It is not a move in our game. But Hansen's point is this: even if we think it

[352] Cf. Kauppinen (2007, 106-7; 109-111).

is safe to assume that our rules are shared by some people, it is still an empirical matter. We are, therefore, mistaken if we think it can be settled by *a priori* reflection alone.[353]

The power of this objection depends on how we answer the following question: is there any obligation for the linguistic philosopher to think the *a priori* method is sufficient to discover such facts? On the one hand, the assumption that *a priori* reflection is supposed to determine the extent to which our rules are shared seems to be Hansen's. Linguistic philosophers map the boundaries of different concepts (often the concepts they themselves employ), but their method leaves open the question Hansen demands they answer. On the other hand, Hansen's assumption is hardly unreasonable. After all, philosophers like Cavell *do* make claims about what *we* say. If knowing what we say is really an empirical matter, then we might be left with the impression that Cavell uses philosophical reflection to deliver conclusions he is not entitled to establish via those means.

But even this problem is not necessarily deep. Not every empirical claim is arrived at via a rigorous scientific process. Nor does every claim require testing. Much of our knowledge is the result of cumulative, everyday experiences. If asked to put it through a process of rigorous testing, we would likely think it a waste of time. And so might we say that of Hansen's challenge. That we form part of a linguistic community is obvious: we go to the same schools, watch the same television programmes, read the same newspapers, listen to the same radio stations, converse with one another without difficulty. All of this is evidence of our sharing a language. Indeed, I have already suggested that consulting how others use words can provide helpful reminders to us when we are trying to provide meaning explanations. Hence, one might think we hardly move from *I* to *We* at all, for the ability we reflect upon is by its very nature social. The thought being that the ability I reflect on, and the conceptual connections I draw, are not merely mine but *ours*.

It is, of course, open to Hansen to reject this move. To remain sceptical of the everyday epistemological claim; to continue to problematise the distinction between *I* and *We*. And in some circumstances, we might deem this appropriate. Successful communication can occur despite many, though subtle, variations in how each of us speak. One might argue that, as philosophers, we should be reluctant to take our linguistic community for granted. This would

[353] 'Alone' may strike us as confused, since *a priori* knowledge typically allows for linguistic knowledge (which we, of course, obtain through experience). Given this is widely taken for granted, I assume 'alone' will not create confusion here. Roughly, *a priori* knowledge can be thought of as knowledge that arises merely from understanding, and thinking about, a proposition.

return us to where we started: Cavell (and philosophers like him) appearing to undermine their own projects by making claims their philosophical methods cannot establish.

7.3 Linguistic proposals

Perhaps the most powerful response available to Cavell is yet to be made, though. In Hansen's eyes, this is certainly the case.[354] He cannot shake off the thought that what we really need, insofar as we are interested in how we use words, is empirical investigation. While he accepts some of the preceding, insofar as he admits 'it is possible to be entitled to make claims about how "we" use certain expressions without engaging in experimental or corpus-based investigations',[355] he still thinks it misses the key claim Cavell wants to make. What Hansen takes from Cavell is that linguistic philosophy's most powerful manoeuvre against the likes of Mates is its appeal to linguistic proposals – or, in modern-day parlance – conceptual engineering.[356] The thought is that philosophers like Cavell can identify important distinctions that we ought to mark with different words. These distinctions might be genuinely novel, or rediscoveries of past distinctions that have since disappeared from our language. Hansen's point, then, is that linguistic philosophers can propose different means of description as being apt or useful, suggesting that we *should* adopt them (not that we already do). This task is safe from the famous criticisms of the empirical semanticists and 'experimental philosophers' because it is not one that empirical research can undertake.[357]

This is certainly something a philosopher can do. Revisions and additions can be, and often are, suggested by philosophers as ways of – in some sense – improving language. Indeed, this may not always be divorceable from discussions about how we do, in fact, use words. One can imagine arriving at a deadlock of the following form, for example. 'But look: the distinction between mistakes and accidents is important. We use it in our day-to-day life to distinguish between different cases'.[358] To which an interlocutor might reply: 'I accept what you say regarding its importance, it just isn't clear to me that we have employed the distinction in the past. But I certainly think we should from now on'. The point being that, for some, the

[354] See Hansen (2017). It is also (briefly) mentioned in Hansen (2014).
[355] Hansen (2017, 812).
[356] On conceptual engineering, see Burgess, Cappelen and Plunkett (2020), Burgess and Plunkett (2013a and b), Cappelen (2018) Eklund (2015), Plunkett (2015), and Plunkett and Sundell (2013).
[357] That conceptual engineering might make room for philosophy in this way is a point made by proponents of that method. See, e.g., Nado (2019) who identifies the challenge from experimental philosophy as one of several problems that can be evaded in this way.
[358] See Austin (1957) for this distinction.

recognition of its importance speaks to the fact it is already employed in their lives, while others might claim it is entirely new to them. Both agree the distinction is worth having. Their disagreement concerns whether this is really a form of conceptual innovation. It may never be settled: they might just agree to go on using the distinction that both find useful.

An example from Wittgenstein's *Philosophical Investigations* might clarify this matter further. Katherine Morris identifies a similar phenomenon in the following passage:[359]

> Suppose someone said, "*All* tools serve to modify something. So, a hammer modifies the position of a nail, a saw the shape of a board, and so on." – And what is modified by a rule, a glue-pot and nails? – "Our knowledge of a thing's length, the temperature of the glue, and the solidity of a box." (PI, §14)

Morris plausibly suggests that the examples raised in the middle of the quotation might be enough to convince an interlocutor to withdraw the initial statement. They might recognise they were working with too small a range of cases. But it is also possible that they respond as Wittgenstein's interlocutor does. What then? We might appeal to ordinary use and suggest that 'modify' is being stretched beyond its usual bounds. But is there any compulsion for the interlocutor to accept this? Even if ordinary use is clear on this point (which is surely a cause for further debate),[360] 'why *should* the interlocutor confine himself to our ordinary use?'[361] Why can't the interlocutor appeal to his own idiosyncrasies? Indeed, this possibility seems to be precisely what Wittgenstein has in mind for them. For the final question raised in §14 does not imply that the interlocutor has made a narrow linguistic error but instead asks whether there is any point in speaking in this manner:

> — Would anything be gained by this assimilation of expressions? – (*Ibid*)

In other words, is this a particularly helpful way of going on? Might we not miss important distinctions should we accept this way of speaking? The interlocutor's intransigence has forced us onto new terrain, the kind Hansen thinks is Cavell's most promising answer to critics of armchair philosophy. Mulhall suggests a criticism one might make of the interlocutor's suggestion:

---

[359] Morris (2019). She uses it to adjudicate over the disagreement that broke out between Baker and Hacker with respect to the nature of philosophy and Wittgenstein's thought after many years of collaboration. I should not be read as endorsing Morris's conclusions about their dispute.
[360] One might think it obvious that we can modify our knowledge, but that 'modify' as used in relation to hammers and saws talks to something more specific (that does not apply to knowledge). Cf. Mulhall in the passage below.
[361] Morris (2019, 123).

> The problem is that each step we take in extending the reach of 'modification' to accommodate apparently resistant cases forces us to attenuate the meaning it has when we apply it specifically to hammers and saws. If temperature and a mental state are to qualify as possible objects of modification just as easily as the position of a nail or the shape of a board, our definition of what counts as an act of modification will become increasingly loose and baggy, and so calling something a 'modifier' will tell us less and less about how it actually works.[362] (Mulhall 2018)

This discussion takes us closer, I think, to Cavell's best defence against his critics, but it is one that neither Sandis nor Hansen really see (or, at any rate, suggest). Hansen is right to recognise the importance of proposals in Cavell's thinking, but he subverts its significance. Not because philosophers are unentitled to suggest new ways of going on, but because proposals of a different kind answer Cavell's critics more directly than Hansen's suggestion ever could. Whereas Hansen locates the critical issue in how we might respond to our intransigent interlocutor (by defending the value of linguistic tools they lack), it is actually manifest in the prior move the interlocutor is entitled to make which might elicit such a response. Let us finish, then, by considering a different form of proposal that occupies a central role in Cavell's thinking.

Who do I speak for when I make claims about the ways in which *we* make sense? Who am I making claims over when I suggest a combination of words is excluded from *our* language? But even this might strike us as a strange way of putting it. It casts the exercise of clarifying our language, determining what makes sense, within a kind of headmaster-pupil context. There is, however, no reason for us to conceive of it like this. Why assume from the outset that I can identify those I'll be speaking for? Is this not something to be worked out in the process of clarification? To take a trivial example: suppose my grandmother tells me that Osborne has come out and 'refuted those claims'. My grandmother isn't much of a fan of Osborne, so it would be some surprise were she to mean what I think she has said. So, I ask her to clarify what Osborne has said and how that proves his case. She tells me she said nothing about him proving his case. Do I reply by saying that she has failed to make sense, for refutation requires proof? Probably not. And even if I did, she would likely respond 'well that wasn't what I meant'. What does this reveal? It reveals that, prior to working through a process of clarification, there is no guarantee we are using our words in the same way. I cannot bind a certain group of people and say: 'this is how you make sense' and go on speaking for them. No – in the process of 'speaking for them', I help determine exactly who I am speaking for. One is entitled at any point to say 'that's not what

[362] Linguistic criticism of this kind can also be found in Cavell's writing (1958 [2002], e.g. 33-4).

I meant', but then one is expected to explain what one did mean. To give sense to a particular usage. This understanding of 'we' is what Mulhall calls 'Cavell's lifelong characterisation of [the] philosophical "we" as entering a claim to community'.[363]

> "But such claims as: 'We say. . .', 'We are not going to call . . . ,' and so forth, are not merely claims about what I say and mean and do, but about what others say and mean and do as well. And how can I speak for others on the basis of knowledge about myself?" The question is: Why are some claims about myself expressed in the form "We . . ."? About what can I speak for others on the basis of what I have learned about myself? […] Then suppose it is asked: "But how do I know others speak as I do?" About some things I know they do not; I have some knowledge of my idiosyncrasy. But if the question means "How do I know at all that others speak as I do?" then the answer is, I do not. I may find out that the most common concept is not used by us in the same way. And one of Wittgenstein's questions is: What would it be like to find this out? (Cavell 1962 [2002], 62)

The philosophical 'we' invites others to make perspicuous their own ways of making sense. It is not a claim to authority. This does not prevent us from acknowledging the 'division of linguistic labour', as Putnam calls it.[364] That we rely on other speakers' expertise and defer to it in the relevant cases is a feature of the linguistic practice we participate in and reflects the relative grasps we have of certain concepts. In deferring, we accept and follow the rules of those experts.[365] Nor does this conception of 'we' remove the possibility of our revising rule-descriptions in light of what others say. After all, it might turn out that we (they and I) do speak as they say, and I was simply mistaken in my previous description. But what it does leave open is the possibility of genuine divergence, meaning 'who exactly "we" may be is itself at issue in every philosophical exchange'.[366] In using 'we', the speaker may assume a certain community as their

---

[363] Mulhall (2009a, 401). See also Mulhall (2015).

[364] Putnam (1973). A classic example is owed to Burge (1979). Bert seeks a doctor complaining that he has arthritis in his thigh. In this case, it is natural to expect the patient to defer to the doctor. This is not only because of the doctor's expertise. The patient wants help and is clearly trying to pick out a condition recognised by the wider community.

[365] Cf. Lassiter (2008) and Ludlow (2014, 83-4). Note that Jackman (2001) cites the division of linguistic labour as his main complaint against this 'individualistic' view of language, but it is not, I argue, something his opponent need sacrifice. Jackman pays little attention to the distinction between use and usage. See also Jackman (2005) in this regard. It might be that his conflation of use and usage prevents him from recognising that acts of deference can reconcile 'individualism' with language's social character. For if usage were correct use, then we might be left having to think only in terms of expert-concepts and folk-concepts, rather than recognising the latter as sometimes simplified, sometimes misunderstood, versions of the former – where those of us with partial understandings of certain concepts defer to the relevant experts. Such behaviour makes plain that correct use is regulated by the community of experts, and that we are not ourselves masters of those concepts.

[366] Mulhall (2015, 154). Gustafsson (2005, 367) also appears to entertain this possibility in the context of Cavell's philosophy.

own, but it is also an invitation which implicitly asks: 'who is willing to travel with me?'. It is not an order to travel with them.[367]

The above understanding of the philosophical 'we' might suggest that any results Wittgenstein arrives at, if one can call them that, are necessarily limited, for he can give no guarantee that his interlocutors' answers would be someone else's and nor can he guarantee his interlocutors' mistakes are typical of his readers. But this is little more than an acknowledgement that his results are definitive of a *particular* way of thinking. Might it help explain the form his writing takes and the role of his interlocutors? It certainly suggests a demand on those engaging with philosophical problems to recognise the scope for divergent understandings of shared concepts.[368] That Wittgenstein thought the main contribution of his work was the development of a philosophical method is therefore perhaps no accident. For however widely the conclusions drawn from his investigations applied at the time, there could be no guarantee they would all have lasting relevance. Not because the results of his method are contingent, but because the practices he investigated may later change, or even be given up altogether.

This point should not be overstated, however. Many practices of the kind we engage in have been with us for millennia and lots of the concepts Wittgenstein explored are so central to our lives that it is difficult to imagine ourselves without them. But we should at least be mindful of the possibility of novel concepts, and changes made to existing ones, that might conjure up new, though related, philosophical problems. Moreover, individual speakers can make idiosyncratic mistakes, the rectifying of which may bear little consequence for anybody else. Perhaps this is one sense in which the comparison between the resolution of philosophical problems and therapy might strike us as illuminating.[369]

8. Conclusion

[367] Sandis (2019) charts the use of 'we' throughout Wittgenstein. While much of what he says there may be correct, I would stress the importance of invoking 'we' as a claim to community (one to be worked out in the process of philosophising) for Wittgenstein's method, even if we may – for the most part – have a general sense of who 'we' are.

[368] Something we already have a nose for insofar as we recognise that one abundant source of philosophical problems is the conflation or equivocation of different concepts.

[369] Cf. PI §133 and §235. This is another area of controversy within Wittgenstein scholarship. For instance, Hacker (2007) downplays the analogy with therapy and criticises Baker's (2004) later work for making too much of it. I should be clear that just because I think it necessary to entertain the possibility of idiosyncratic mistakes and ways of speaking does not mean I suppose they are ubiquitous and hence characterise the nature of philosophy. Philosophers, Wittgenstein included, presumably prioritise the rectification of common mistakes. As Hacker notes, when Wittgenstein discussed Hardy's claim that 'a reality corresponds to mathematical propositions', he added that the 'fact that he said it does not matter; what is important is that it is a thing which lots of people would like to say' (LFM 239).

By recognising the possibility of idiosyncratic ways of speaking, Cavell answers his critics. Philosophical reflection is not, by itself, capable of delivering knowledge of others. About that, his critics are correct. But it was never supposed to. It is one thing to understand a system of norms and what follows from them, quite another to determine who uses those norms. The former is a conceptual, or logical, exercise. The latter is clearly an empirical one.

Philosophical dialogue is an exercise fit to contribute to both our conceptual understanding and empirical knowledge. We stand to learn from such discussions not merely what follows from a particular set of norms, but who follows which set. In the process of resolving disagreements, we may discover that we differ in how we speak. An argument regarding which is the better way of speaking, or thinking, might ensue, but it is not compulsory. Idiosyncrasies can be tolerated and need not lead to confusion once they have been clarified. On certain topics, we might feel compelled to try to convince others of our way of thinking. This could be due to the theoretical advantages of certain concepts, the ethical reasons we have for our preferences, or something else.[370]

None of this is to deny that we are often explicit in justifying our linguistic corrections by appeal to a generally shared language ('You can speak like that if you want, but it isn't what is usually meant by that term'). After all, there are clear advantages to speaking as others do. Furthermore, children may not be given the opportunity to expound their idiosyncratic meanings. This might be because we doubt they have any. They are simply making mistakes as they try to grapple with our meanings. But the more salient reason is that the aim of their education is to introduce them to our language and way of life. We are initiating them into our community and so they need to understand how we speak. This is different from cases involving mature language-users. Take one who, for example, refuses to acquiesce to experts with respect to what constitutes a fish. Suppose they include both whales and dolphins under their classification. If our deviant speaker really has no interest in following us or acquiescing to those with the relevant technical knowledge (and consistently applies their term according to their own criteria), we may eventually have no option but to accept they use a different concept.[371]

Thinking about how we use words is a matter of reflecting on the ability we have and the norms that constitute it: what we would say in different circumstances and why we would say it. But when *I* make a claim about how *we* speak, it cannot be taken for granted that I know exactly who I am speaking for. That is something we learn in conversation with others. Perhaps, then,

[370] See Part III for further discussion.
[371] Donnellan (1962, 650) makes a similar point using this example.

we should say that philosophical reflection can tell us what we mean, but not – by its *a priori* methods alone – who we are.

Part III

The argument of Part I was that necessary truths are insubstantial and best understood as being conceptual in nature. The second part defended the distinction between empirical and conceptual propositions and argued that it can be maintained despite the presence of indeterminacy in particular cases. The third and final part now asks what the relation is between the world and conceptual propositions. If it is right that conceptual propositions are distinct from those that describe the world, how should the relation between those conceptual propositions (and hence our conceptual resources) and reality be understood? The first chapter deals with this question directly. The second considers what that answer means for the relation between 'necessary truths' and worldly facts and defends the normative account of necessity against an interpretation that transforms it into a kind of Kantian linguistic idealism. This interpretation is the result of misconstruing the relation between the facts and conceptual truths.

This is clearly important for the sake of completing the account. If I am denying that conceptual truths answer to the world, I owe an explanation of what the relation between them and the world is. Aside from this, certain worries arise from seeming to divorce conceptual truths from worldly facts, for our concepts and the facts seem to be very well aligned. Moreover, we seem to be justified in using the concepts we do, so what accounts for that justification if not that they – in some sense – mirror the world? Thomasson outlines three related concerns in this regard.

> [A]bandoning a metaphysical approach to conceptual choice leads to three interrelated worries: (1) That we will be unable to account for intuitions of structure; (2) That we will have to treat conceptual choices as merely arbitrary, not worldly; and (3) That we will be unable to critique conceptual choices. (Thomasson 2020b, 439-40)

Thomasson thinks that none hold of the normative account. I agree and in what follows I try to articulate a picture wherein our concepts are intertwined with the conditions they're used in but do not answer to them in the way that descriptions do.

(9) Concepts and Facts

In chapter seven, we saw how conceptual change could occur throughout scientific advance without compromising popular conceptions of science. Scientific discourse is just one of many, though it is especially important insofar as it might have proved awkward for a linguistic account. One potential difficulty, to recall, was that we are developing tools to help us best describe and explain reality, and hence it is easy to mistake this fact for thinking that the relevant norms we institute are themselves answerable to the world in the same way that the descriptions which use them are.

The task now is, in a sense, to generalise that discussion – to study how our lives and general facts about the world relate to our conceptual scheme. The aim is to clarify the relation between worldly facts and our linguistic resources. There are several ways one might go about this. My method is to start by looking at conceptual disagreement. What kinds of considerations are there when deciding between different concepts? What role do facts play? We will consider different types of disagreements, the reasons available to us when involved in them, and what is at stake in the process of resolving them.

1. Disagreement

Disagreement comes in many forms. The most straightforward is when we disagree about something within a language-game. A standard example would be disagreeing over the facts. We might see an advert that you claim contains a leopard while I say it is a jaguar. We might disagree about what should be done to achieve a certain aim. We can disagree over how something tastes – whether it needs any salt, sugar or additional spices adding. All these disagreements occur inside a language-game, which is to say the norms that constitute those games are neither in dispute nor up for grabs. We are speaking a shared language, rather than trying to establish one.

But sometimes when we disagree about whether, for example, a given object qualifies as a certain type, it is not because we disagree about the facts. Rather, what we disagree about is what it takes, sometimes what it *should* take, to qualify as that type. We can therefore distinguish between two kinds of disagreements that concern linguistic norms. Some will concern matters of fact – facts about our semantic histories and the structure of the language we speak. Others will not be matters of fact at all. They will involve practical questions about which meanings we want to go forward with. We can summarise this by saying there are:

i. Disagreements over what reality is like;
ii. Disagreements about what to do;
iii. Disagreements over what our words mean; and
iv. Disagreements about how we should speak and make sense of things.[372]

Disagreements of the third type can be broken down into further kinds. First, we might simply disagree about the meaning of a particular word. For instance, someone might confuse the meanings of 'disingenuous' and 'ingenious'. Here we might say there is no conceptual mistake, only a lexical one. They understand the relevant concepts but confuse the correct words. This might be typical of someone learning a second language. We might say someone makes a *conceptual* mistake when they get the norms wrong in some way. For instance, by claiming 'it isn't a mammal if it lays eggs', only to be reminded that both the duck-billed platypus and echidna are egg-laying mammals. Now, these types of disagreements will typically be settled. Most of us accept that we belong to a linguistic community whereby the meaning of 'mammal' does include those animals (even if we are unaware of that fact). Likewise, a dictionary definition will usually suffice to convince us that we have mistaken 'disingenuous' for 'ingenious'.

In other cases, however, words may either admit of several different senses, or we may simply use a word idiosyncratically (and know this fact about ourselves). In such cases, these inconsistencies do not rise to the level of disputes. There is no contest. It is just that our uses of the relevant terms do not coincide. Brendan Jackson gives the example of 'metaphysics'. In (American) English, it can 'cover the study of various supernatural phenomena, such as out-of-body experiences, past lives, and the use of crystals to "channel energies."'[373] So, an analytic philosopher might criticise the size of the 'metaphysics' section of a library, while someone else might claim the library contains no books on metaphysics at all. Here they do not disagree about the facts and need not dispute that both uses of 'metaphysics' are permissible. They are merely using their words in legitimate, but different, ways.

On occasion, disagreements of this kind might lead to the fourth. We might want to contest the other person's standards, because we think ours are better for some purpose. Peter Ludlow presents the example of a radio programme debating whether it was legitimate for a magazine to list a racehorse as one of the twentieth century's best athletes.[374] Whatever we think 'athlete' means (it seems to me to exclude horses), the discussion could easily stray into the

[372] They are not always easily distinguishable, as Part II should have made clear.
[373] Jackson (2014, 33).
[374] Ludlow (2008).

question of how appropriate it is to amend our concept to include racehorses. The magazine presumably knew it would be a controversial choice, and it was – in effect – challenging our ordinary conception of what an athlete can be.[375] Another example might include the following. You and I are setting up a restaurant and we need to decide at what point we mark 'hot' next to certain dishes on the menu. That is, how spicy a dish can be before being listed as such.[376] I might argue that we should mirror the average tolerance of a typical diner, while you suggest that, given the style of restaurant we are opening, the threshold should be higher, as diners will expect 'hot' to signify as much.

This fourth kind of disagreement is the one we are most interested in. A recurring topic throughout this thesis has been the potential for novel linguistic proposals. It is especially pertinent in this context. The question regarding how we might decide which standards to employ requires us to think about the kinds of considerations that might affect our language. What are the conditions that affect our decisions, that determine the kinds of concepts we have? Of course, much of our language we did not 'decide' upon. It was inherited. But it was devised by a community. Not through the writing of a rulebook initially (for no such book could be written without a language to write it in), but (presumably) through its consistent use within their lives such that norms became established. Moreover, language is under constant (both conscious and unconscious) development. The kinds of reasons we have for preferring one mode of speaking over another play a role in the development and maintenance of our language. The general question we are interested in is the following: what are the conditions that might determine the kind of conceptual scheme we have? How do the facts (broadly construed) impact our language? What is the relation between them? In thinking about how language is used and what it is used for, we should begin to establish a picture of the interaction between the facts and conceptual norms. One that is consistent with thinking those norms are not 'made true' by the facts.

2. Conditions that shape how we make sense

[375] Deliberate transgressions like this are explored in Sterken (2020).
[376] Plunkett (2015, 839) suggests an analogous example concerning a mini-golf course and the level of difficulty assigned to different holes. Plunkett and Sundell (2013: 14−15) use 'spicy' as an example, though not in the context of opening a restaurant.

> It is a fact of experience that human beings alter their concepts, exchange them for others when they learn new facts; when in this way what was formerly important to them becomes unimportant, and *vice versa*. (*Z*, §352)

> … but has nature nothing to say here? Indeed she has – but she makes herself audible in another way.
>
> "You'll surely run up against existence and non-existence somewhere!" But that means against *facts*, not concepts.[377] (*Z*, §364)

Let us adapt our kitchen example and assume we have established a standard already. Suppose now you enter that kitchen as a chef with a very high threshold for spice and a correspondingly high standard for the term 'hot'. If your customers do not share it, you might be required to adjust your (linguistic) standard (at least while you are in the kitchen). There may be an accepted threshold of when something is 'hot', which is adjusted to the preferences of the diners, that you need to learn to understand your instructions. Otherwise, you risk hearing an instruction that has not been given and ruining the diner's experience as a result.

The kitchen case illustrates some of the conditions that shape the concepts we use. We can list some of those general facts. First, we have senses of taste and different foods have different tastes. Second, some foods are hot. Third, tolerance and enjoyment of spice varies person to person. Fourth, we sometimes cook together and for each other. Fifth, we usually want our guests to enjoy their food. Sixth, there are economic incentives that make it desirable that paying guests enjoy their food. So far, these facts are very general, but they help explain why we have the concept in the first place and why it might be important for some of us to share the same standards. We can also add a seventh, more specific fact: guests at this restaurant tend to have a certain tolerance of spice. This condition, in our example, provides the reason for setting the standard where we do. It is useful to have these standards in place because it is in our interests to ensure that the diners enjoy their food.

But that the kitchen's use of 'hot' corresponds with its guests' tolerance does not make it right. What those tastes are is, of course, a matter of fact. But even if the standards have been set with those tastes in mind, the concept is not correct because those standards match those tastes, for conceptual norms do not answer to those tastes. They are not an attempt to describe the facts but constitute a mode of describing them.[378] Of course, there are common standards for the

[377] See Hacker (1997, 190-1) for a discussion of this point.
[378] Cf. Mulhall's (2009b, 153) gloss on Hacker's interpretation of Wittgenstein.

English word 'hot', and we would expect the restaurant to be aligned with them, but the word is not sufficiently precise to determine the kitchen's decision. Hence restaurants can have (somewhat) idiosyncratic standards against which their chefs' use of the word, for example, could be judged. The main point to make is that there is a significant difference between saying 'this is the point at which something will count as "hot"' and reporting the point at which diners say a dish is hot. This is not to say there is no rationale for these standards. It is just to say that while we might measure our success in developing the concept against those facts, it does not follow that those norms are answerable to them. Those standards are one set of many and that they are preferred is a matter of how the facts interact with the kitchen's interests. Moreover, the kitchen could just have easily opted for a different standard and changed how they talk about their cooking and kept it to their diners' tastes. In a similar vein, while a rule change in sport might be designed to make it a more enjoyable watch, or safer to play, it does not follow that those rules are true or false just in case they achieve those goals. Statements of those rules are counted true when they reflect the rules of the game. But we might say those rules, or rule-changes, are *vindicated* when those stated aims are achieved.

This is not altogether different from the scientific case, where certain discoveries can lead us to change our conceptual tools, and various empirical facts may be hardened into rules.[379] Properties that are found to be typical of a given phenomenon can be transformed into criteria, altering what it is to be a phenomenon of that kind (though the extension of the relevant concept may barely change). This is plausibly what occurred when we defined acids as proton donors.[380] These developments do not render earlier rules incorrect. Rather, they tend to indicate that earlier rules determined less efficacious concepts.

This, as Thomasson argues, helps account for certain 'intuitions' regarding metaphysical structure.[381] One reason our current chemical concepts are better when compared to, say, an imaginary set that enforce geographical restrictions (for example, 'lithium-in-Cameroon') is that geographical constraints are not, in and of themselves, especially relevant to scientific explanations or predictions. Hence, such concepts would be of little use in chemical theory. Likewise, the concepts 'fish' and 'mammal' are better than 'sea creature' and 'land creature' insofar as they are designed to figure in biological explanations. These concepts are better

[379] Cf. RFM VI §22-3; OC §321.
[380] Hacker (1996a, 215).
[381] Thomasson (2020b).

indications of internal structures, disease susceptibility, and behaviour.[382] We can say they are more natural by appealing to *empirical* facts. Metaphysical super-facts are entirely unnecessary.

This is something Thomasson is keen to emphasise. Her child has a nut allergy, so it is a matter of life and death 'whether something is biologically a tree nut or is something *called* a "nut"'.[383] Hence, the concept 'tree nut' is especially useful because it picks out those items that could send her child into anaphylaxis. Of course, there may be other reasons why it is useful, but this one has special value for Thomasson.

> The important thing to note here is that there is a way to justify the claim that one concept is better than another, for worldly reasons, but that our choice of concepts is vindicated *empirically*, given our shared purposes—it does not require additional *metaphysical* vindication. (Thomasson 2020b, 450-1)

Wittgenstein suggests something like this applies even to elementary mathematics. We can imagine situations where our arithmetic has few applications in the world, because the world behaves in unusual and irregular ways such that our method for transforming quantities rarely applies to it. If objects tended to change shape, multiply, and vanish at random, the use of our basic mathematics would be significantly curtailed. It would be unreasonable, for example, to infer from there being six apples to a box and ten boxes, that there are sixty apples. This is not a matter of disputing that six multiplied by ten is sixty. That much remains necessary on Wittgenstein's view, as do other calculations. After all, we could plausibly describe the scenarios in our usual terms, even if those descriptions tended to be not very useful. We might notice that three had disappeared by counting fifty-seven apples, for instance. But the point is that if this too were likely to change, then even the act of counting those apples would be largely pointless, or, at least, wouldn't have the purpose it usually does.[384] It might not even tell us how many there are: the number may change over the course of counting. In other words, our way of representing things is not as useful in this alternative world. There is a certain stability necessary for the sums to be worth doing. Perhaps if objects behaved in strange, but predictable ways, we might have a different arithmetic that mirrored that behaviour. This possibility Wittgenstein takes very seriously.[385] It is worth noticing that there are units we use now that do not behave according to basic arithmetic. As Schroeder notes, two teardrops added together make one

[382] *Ibid,* 450.
[383] *Ibid.*
[384] Cf. Waismann (1953, 74-5).
[385] Cf. RFM I §38; §140).

teardrop and one quart water added to one quart alcohol makes only 1.8 quarts of vodka.[386] The point is not to deny that we have explanations for this, but to suggest that if these anomalies were to generalise, what we take as basic arithmetic may resemble something quite alien to us.[387]

There is a myriad of reasons why a given concept might prove less useful than an alternative. Fundamental to this point is that we can evaluate new concepts by asking whether they succeed in performing the function we intend for them.[388] Some might be designed to replace existing ones, others may be invented not as replacements but additions.[389] Old concepts may be repurposed, or even abandoned altogether, following a critique of their former use or the practice they played a role in.[390] The ways in which concepts may be deficient can help shed light on the intimate connection between our language and the facts. Herman Cappelen makes a start on a taxonomy of sorts for conceptual deficiency.[391] Some examples include semantic failures (nonsense, reference failure, vagueness in circumstances where precision is required), cognitive defects (where certain concepts are empirically associated with flawed reasoning), social/political failures (concepts which promote projects we think are morally reprehensible), and detrimental effects on theorising. These are neither exhaustive nor mutually exclusive categories. But they do offer a view of conceptual deficiency in which facts play a role. Whether a concept is helpful in theorising might depend on the causal significance of the distinction the concept draws. Likewise, if a word fails to refer, half of the story is the non-existence of the phenomenon in question. Social and political failures similarly rest on social realities, and no doubt the histories of different societies and their linguistic tools. This begins to demonstrate the intimate connection our concepts have with the facts, as tools for us to use in the world. Language is something we use, and different language-games generally have some purpose and form part of a wider activity or practice.

What we have, then, is a range of facts that help determine our use of words and concepts, where this can be explained with reference to the lives in which they are used. These include facts about the world and facts about us: our cognitive and sensory capacities, interests, knowledge, and values. Some of these facts will be very general, and these tend to explain why we have a certain category of concept, while more specific facts may explain why we employ the particular concept we do. This is true as much when comparing options within the same

[386] Schroeder (2015, 125).
[387] For more on this, see Bangu (2019) and Schroeder (2021).
[388] Simion and Kelp (2020, 986).
[389] See *Ibid* for a convincing defence of this point.
[390] See Nado (2019, 1519).
[391] Cappelen (2018).

language as it is between them. We have cross-cutting categories, appropriate in different circumstances. For example, 'tree' need not be a taxonomic group to prove useful to landscape gardeners.[392]

3. Life and language

These observations suggest a connection that is – in a sense – deeper than it first appears. For without some of those facts being in place, the concepts conditioned by them would not be in use. Circumstances call for certain concepts. Some functions are not there to be fulfilled in different contexts and – as we saw with open texture – there are contexts where it just isn't clear what counts as falling under a concept that is prepared for altogether different circumstances.[393] Similarly, if we were substantially different from what we are, certain concepts might be unusable. This is explored by Wittgenstein in *Philosophical Investigations* and several of his other works.

> It is only in normal cases that the use of a word is clearly laid out in advance for us; we know, are in no doubt, what we have to say in this or that case. The more abnormal the case, the more doubtful it becomes what we are to say. And if things were quite different from what they actually are — if there were, for instance, no characteristic expression of pain, of fear, of joy; if rule became exception, and exception rule; or if both became phenomena of roughly equal frequency — our normal language-games would thereby lose their point. – The procedure of putting a lump of cheese on a balance and fixing the price by the turn of the scale would lose its point if it frequently happened that such lumps suddenly grew or shrank with no obvious cause.[394] (PI, §142)

We can think of the final example from the perspective of the consumer and producer of a particular type of cheese. For the consumer, the price is supposed to reflect the amount of cheese the customer will have to consume. The more they pay, the more they get. For the producer, price differentials of the same cheese will typically reflect the relative cost of production (larger blocks contain a greater quantity of ingredients). But if the final product is likely to change in size at random, then fixing the price of cheese by its weight at a given time

[392] Cf. Dupré (1981, 80-2; also 1993).
[393] Z, §350.
[394] See also PI §200; Z §350-351; PPF §346.

will not reliably reflect either of these things. The point of fixing the price of cheese by the turn of the scale is thereby challenged. The concept no longer has the same use or significance.[395]

Thus far we have stressed one side of the interplay between life and language, that being the role our circumstances have on the concepts we use. But concepts have a kind of enabling feature: they grant us access to the facts and to different ways of living and acting. To bring this out, recall a distinction I drew in the first section between lexical and conceptual mistakes. The difference between getting the words wrong (confusing 'disingenuous' with 'ingenious') and getting the conceptual norms themselves wrong (claiming something cannot be a mammal if it lays eggs). In the context of a language undergoing alterations, the difference is between changing the words we use (which might involve creating or abolishing a use for a particular word) versus changing the terms in which we think (how we make sense).[396]

This distinction is made vivid by Oswald Hanfling who shows that disusing a word is not tantamount to abandoning the concept thereby expressed.[397] The first example he suggests is the distinction between refutation and rejection. The two uses have been eroded in recent years, yet this hardly threatens the conceptual distinction that was once marked by those words. It would be of much greater significance were we really giving up on the possibility of disproving something, and instead had room only for people's performative assertions to the contrary. Indeed, a life lacking a system of proof would be dramatically different from the one we live. Would the activity formally known as 'proof' even take place? Something might mimic it, but the very notion of proof would have disappeared (that being the idea that something is established as true by means of evidence or argument). The consequences of this mimicry could only be our rejection of the relevant claim with no indication that the matter had been (nor could be) settled. Hence, we would lose the sense that someone who comes across the proof *ought* to reject the relevant claim. But, of course, none of these things follow from the blurring of words. The act of refutation continues to play an important role in our lives. What refutation consists in can be elucidated by studying what it is we do in refuting something. It can be ascertained by studying

[395] Mulhall (2001, 99-105, 122-135; 2009b, 162-5), following Diamond (1989), develops this passage in a striking fashion – going as far as entertaining the possibility that under these circumstances one wouldn't even be fixing the price of cheese. Diamond and Mulhall aim to use their readings of similar passages to resolve the debate between Baker and Hacker (1990) and Malcolm (1989) on the communal view of rule-following.
[396] The latter is more interesting, but it is worth noting that the former is no less conditioned by the facts and can still be evaluated. Some term might have problematic connotations given its history, for example.
[397] Hanfling (2000, 60-65).

the role it plays in the lives of those in possession of the concept.[398] And whether a different community has the concept is a question of what they do, not which words they use.[399]

Another example of Hanfling's he takes from Cook, who adapted an example discussed by Austin and Cavell.[400] We are to imagine two separate stories. In one, a baker's assistant arrives and turns the ovens on at the usual time despite having been asked to turn them on later. In another, the assistant remembers to turn them on later, as planned, but when they go to switch the lights on, their fumbling ends up switching the ovens on as well. The claim is that in the first example the assistant turns the ovens on *automatically*, while in the second they do so *inadvertently*. Important for our purposes is the fact that we could recognise this distinction in the lives of the baker and their assistants even if they used 'automatically' and 'inadvertently' interchangeably. A distinction between the relevant actions could still be acknowledged by them. For instance, in the first story, the baker might suggest that a sign to remind the assistant not to turn the ovens on would have prevented the mistake. But this response would be nonsensical in the second case. Perhaps instead the suggestion would be to tell the assistant to be more careful, or to check that they haven't switched the ovens on without noticing. What they would say if asked why they wouldn't make the same sign in the second case might show they understand those cases in the same terms we do. They might say: 'the assistant neither meant to do it, nor knew that they had, so that sign wouldn't prevent the same mistake'. In colloquial terms, they might say in the first case one's mind slips, while in the second it is one's hand, and these call for different responses. A community that responds like this recognises the distinction, whatever their terminology.

Concepts shape how we understand ourselves, each other, and the world around us. While language may be thought of as a largely abstract practice that allows us to discuss and represent various aspects of life, the truth is that our concepts are made manifest throughout it. They are not merely used to talk about our lives but play an important role in how we live. 'They are the expression of our interest and direct our interest.'[401] Certain interests can account for the concepts we have (which we largely inherit from others) but once we have them, they structure how we live and how we interpret the world. Hence the need (or want), sometimes, for conceptual change. Related to these insights is the claim that to understand our actions, to see what it is we are trying to do, can require an understanding of the way we make sense of ourselves.[402] Think, for example, of how prayer, crowd reactions to sporting fixtures (or indeed

[398] Cf. Glock (2017, 98).
[399] Cf. Canfield (2009, 112-3).
[400] Cook (1982, 219).
[401] PI §570. Cf. Austin (1957, 130).
[402] Cf. Janik and Toulmin (1973, 228).

the fixtures themselves), and a King's coronation might look to those unfamiliar with them.[403] One has to recognise the role those actions are given to understand what is going on.

To provide an explanation of a concept is, in effect, to say how it is used in the lives of those who use it. While our norms apply in typical contextual conditions, it is far from guaranteed that they can be meaningfully applied in the same way outside of them.[404] There may be no place for our concepts in vastly different circumstances. What we can do in one context may not be replicable in another. Conversely, the concepts we have will inform what we do and so impact our circumstances. In exploring this, Wittgenstein draws our attention to the intimate connection between the lives we lead and the concepts we use. He can be seen throughout his philosophy as trying to grapple with exactly how best to articulate the relation between life and language. Here is one such attempt:

> Would it be correct to say our concepts reflect our life?
>
> They stand in the middle of it. (RC III, §302)

4. Mere semantics

This should put paid to the idea that investigating concepts (or changing them) is a trivial exercise. For determining the nature of the distinctions we draw, or arguing for a new set to be drawn, is not reducible to quarrelling over words. To be sure, we can quarrel about words. We can argue that it would be better if we didn't erode different uses of two words for the sake of effective communication. If we already have two expressions signifying different things, it might be unwise to assimilate those expressions. But these matters are of relatively little significance compared to their conceptual counterparts. Conceptual investigations are studies of what it takes to be an object, action, relation (etc.) of a certain type. And, as we have seen, the concepts we employ play a large role in how we judge, speak, think, and act. The categorial structure of our thought has far-reaching effects.

> Arguably, our conceptual repertoire determines not only what beliefs we can have but also what hypotheses we can entertain, what desires we can form, what plans we can make on the basis of such mental states, and accordingly constrains what we can hope to accomplish in the world. Representation enables action, from the most sophisticated scientific research, to the most mundane household task. It influences our options within

[403] Cf. RFM I §153.
[404] Cf. Z §350.

> social/political institutions and even helps determine which institutions are so much as thinkable. Our social roles, in turn, help determine what kinds of people we can be, what sorts of lives we can lead. (Burgess and Plunkett 2013a, 1096-7)

The extent to which we can *choose* our conceptual scheme is clearly limited. There will be certain distinctions that are unavailable to us due to our capacities (or lack thereof). It is also difficult to see how, for example, we could abandon notions of proof and evidence, whether that is due to a form of pathway dependence or something else. Logical and mathematical reasoning (at least as we know it) would be absent. The notion of matters having been settled, that there are certain things we ought to believe would be compromised. It is to imagine an entirely different mode of living, one where much of what we take for granted would be displaced and many of the activities we otherwise engage in foreclosed. Nevertheless, there are clear examples of choices that we can make.

Thus, the claim that conceptual disagreements are *merely* semantic (and, therefore, trivial) should be challenged. It was discussed briefly in the fourth chapter. Here it might be recast. Insofar as these disagreements matter, it might now be said that what is construed as conceptual change (or engineering) is too serious to be a question of mere semantics. Arguments regarding what can or cannot constitute a marriage, for instance, are not best understood as metalinguistic disputes about what 'marriage' should mean. They're about what our institution of marriage should be. This can't be a conceptual matter, they might say, because it is too important. The meanings of words are, relatively speaking, unimportant.

We should, however, argue that, on the contrary, all this shows is that conceptual matters can be important. In as much as it is important to us what constitutes marriage, it is important to us what the word 'marriage' means. While the mere attaching of meanings to sequences of letters is unimportant, which meanings (or concepts) are in-play is not trivial for reasons already discussed. Where concepts (continue to) play an important role in our lives, such as in the case of 'marriage', changing them reflects how we want to live and so goes beyond the mere attaching of meanings to labels. In deciding (or contesting) which of the meanings should be established, we are determining (or debating) which legal and social relationships are to be possible in our community. Legal definitions may have far reaching consequences depending on the legislation. And more informal social norms can have just as powerful an effect, even if how they do so is less explicit. That we are talking about legal and social relationships shapes the arguments we have. For it is intimately connected to the lives that members of our community can lead. It is crucial to note here that the argument is over which concept ought to be adopted, or accepted,

by the wider community. Sub-groups may already employ concepts they wish to be taken up, but their social aims will only be met if they're taken up more widely (or enshrined in law).[405]

This intimate connection between our lives and concepts was also present in the distinctions between 'refutation' and 'rejection', and 'automatic' and 'inadvertent'. For those concepts were, in part, constituted by what follows from the different descriptions: what we can and should do in response to different kinds of actions. Most would agree that our lives would be significantly poorer without them. Knowing the difference between doing something inadvertently and automatically, for example, allows us to tailor more specialised (and suitable) responses to actions of each kind. Having access to these distinctions is not trivial, even if it can be appropriately understood as a semantic, or conceptual, matter.

5. Genealogy

The foregoing has an interesting connection with genealogy. Some disregard genealogical thinking on account of the genetic fallacy. But this is much too quick. While it may be true that a given claim should be judged on its merits rather than its origins, it doesn't follow that its origins are entirely irrelevant. As Matthieu Queloz points out, justifications for certain practices rely on the origin of their construction (think, for example, of procedural justice or religious practices that are justified by reference to their authoritative origins).[406] Genesis and justification come together when something is genetically justified. Moreover, my confidence in a claim very often depends on its source. Genealogy can be epistemically undermining through the suggestion that one believes something for bad (irrelevant, even) reasons. It may also help to undermine speculative epistemic processes, insofar as it can account for the relevant beliefs without them.[407] Conversely, genealogy may show the belief was formed using a process we know to be reliable. Critics of the genetic fallacy are, of course, quite correct that this cannot show whether the claim is true or false.

More importantly for my purposes, the assumption that genealogical thinking (in general) aims at assessing the epistemology or merits of claims is unwarranted. Genealogical thinking can

[405] Cf. Sterken (2020).
[406] Queloz (2021, 215-6).
[407] Nietzsche (1887 [2006]) is interpreted (separately) by Geuss (1994) and Kail (2011) as targeting this type of epistemic disruption with respect to the origins of moral values. For Geuss, it is intended to undermine beliefs about where Christian values come from and in turn undermine the Christian's belief in their values, while Kail thinks revealing the true origin of those values is designed to question our capacity for describing moral reality. See Srinivasan (2019) for more on this.

play a role in establishing the place certain concepts have in our lives: what functions they serve and the interests that may have shaped them. It can also serve to suggest who stands to benefit from their continued usage. Historical genealogy aims to show the true origins of a concept, and perhaps how it has developed and been sustained over time, while hypothetical genealogy imagines circumstances in which a concept might have been adopted and found useful. The two may be combined, as both Queloz and Williams suggest.[408] We might begin with a highly idealised model, only to later elaborate it with genuine historical details. Of course, historical genealogy's main aim might simply be to provide a causal-historical account, but that doesn't prevent it from illuminating (or articulating) conceptual practices.[409] Both historical and fictional genealogies have the potential to reveal the work any given concept does, by interrogating the circumstances in which it might (or has) come into use, and those where it might be (or has been) deemed advantageous.[410]

Philip Pettit excavates a nice example from H.L.A. Hart's *The Concept of Law*.[411] Assume most societies require some accepted norms of behaviour to function. Next, recognise that informal and uncodified norms can lead to several problems, including problematic vagueness with respect to what is required by the norm, lack of flexibility for new circumstances, and inefficient means of settling disputes between victims and perpetrators. A natural response to these problems is to begin to develop something akin to law. What matters here is not that this is how law was brought about, but that it is a genuine possibility which can illuminate what law does: 'it directs us to functions that law is conceptually required to serve.'[412]

> Relying on the genealogy to identify the referent of the concept, it points us towards the sort of property it serves to predicate. And to that extent it achieves at least some of the purposes that analysis and allocation, pursued as such, might have hoped to achieve. It removes any mystery as to how we could get the concept of law going, it makes sense of the role of the concept without debunking it, and it directs us to a plausible property that constitutes its referent: this is what the concept serves to ascribe. (Pettit 2020, 350)

This potential use of genealogy supports the attention Wittgenstein pays to the point, or purpose, of words, rules, and games.[413] Understanding what the point of a set of rules is, or the

[408] Williams (2002); Queloz (2021).
[409] In his review of Queloz's book, Kail (2021) raises doubts as to whether some of those engaged in historical genealogy, that Queloz claims for his tradition, really engage in genealogy for the purpose of conceptual articulation (though this doesn't mean it couldn't be a by-product of their activity).
[410] See Queloz (2021) for a book-length discussion of this – what he calls 'pragmatic genealogy'.
[411] Pettit (2020); Hart (1961).
[412] Pettit (2020, 350).
[413] PI §563-570.

purpose a concept is supposed to fulfil, helps us to discriminate between good and bad analyses of them.[414] A suggested analysis might, for example, render the concept pointless, or inapplicable, in the circumstances it is designed for. As Pettit suggests, scrutinising the function something is conceptually required to serve achieves, if it doesn't itself constitute a form of analysis, some of the results that analysis otherwise aspires to.

Such genealogies may come to undermine our conceptual practices or vindicate them.[415] By exposing those functions and purposes, we can see the extent to which they are met by the existing practice (or might be better met otherwise) and ask whether those functions and purposes still meet our needs or wants. Genealogy may also serve to reveal connections that are not themselves conceptual. It can, of course, do both at the same time. Non-conceptual connections are not constitutive of the discourse. They are parasitic on it. Once some discourse is defined (and it, therefore, has a general function), it can be put to many different ends. Analogously, during a conflict, one can use the concerns of one side to distract from the more serious ones of the other. But it doesn't follow that the former's concerns are illegitimate or a distraction in and of themselves. Rather, it is just the function they are playing in the present context. Such functions can be institutional, systematically reinforced. This distinction may be like that which is sometimes drawn between semantics and pragmatics.

A recent example of genealogy aimed at unearthing such connections might be found in the work of Samuel Moyn.[416] He critically assesses the way human rights discourse has been used to advance (or not) global justice. The point, Amia Srinivasan suggests, is not to argue that this is an essential feature of the discourse itself, but rather to expose the way the discourse has been systematically used to achieve certain ends. Indeed, she suggests this might help explain 'the extraordinary success of human rights discourse.'[417] In other words, not only can genealogy help unearth certain functional requirements of a given concept, but it may also stimulate criticisms of that concept. That is, it might reveal some of the deficiencies we previously discussed – and, perhaps, some virtues as well.[418]

---

[414] Cf. Queloz (2021, 46).

[415] The notion that some genealogies may be 'vindicatory' is taken from Williams (2002). He suggests, for example, that Hume's account of justice may increase one's respect for that concept and related claims 'if one had [previously] suspected that justice had to be a Platonically other-worldly idea if it was anything' (*Ibid*, 36). Think also of Hobbes' state of nature, which might be said to be designed to justify absolute sovereignty. See Lorenzini (2020) for another conception of genealogy, something he calls 'possibilising genealogy'.

[416] Moyn (2010; 2018).

[417] Srinivasan (2019, 142-3).

[418] For connections between conceptual engineering and genealogy, and further discussion about how the former may be supported by the latter: Dutilh Novaes (2020) and Queloz (2021). This seems a natural suggestion, insofar as genealogy might (i) help us understand the concepts we aim to improve, (ii) reveal the purposes and functions they

6. Conclusion

The task for this chapter was to consider how aspects of the world and our lives shape the concepts we use. Given the role concepts play, it is unsurprising that the connection between our lives and those concepts is deep. That one's life shows the concepts one uses. The deep rootedness of our concepts has also helped to clarify the significance of our conceptual practices.

might play (allowing us to consider how those purposes might be better met), and (iii) expose deficiencies (and/or virtues) the concepts may have.

(10) Linguistic Idealism[419]

This chapter builds on the last by returning to where the thesis started. How should we understand the relation between 'necessary truths' and the world, given the relation between the concepts and facts just sketched? This is particularly important because a common misinterpretation of Wittgenstein's later account of necessity, and the kind of account I have thus far proposed, characterises it as a kind of linguistic transcendental idealism. Not only is this a misinterpretation of the account, but it is also a philosophically hopeless position. Thus, to undermine this misinterpretation is to defend the account's picture of the relation between necessary truths and worldly facts.

To understand the Kantian interpretation of the normative account of necessity, it is best to study the authors who interpret Wittgenstein in this way. For that reason, this chapter features some Wittgenstein scholarship. For what I aim to show is that Wittgenstein already has the resources to resist the Kantian interpretation. Indeed, these resources are central to the normative account of necessity and hence any objections based on this interpretation are unsuccessful. Ultimately, however, it would be enough for my purposes could I show that *some* normative accounts can resist the Kantian interpretation.

1. The non-arbitrariness of grammar

> [I]f anyone believes that certain concepts are absolutely the correct ones, and that having different ones would mean not realising something that we realise – then let him imagine certain very general facts of nature to be different from what we are used to, and the formation of concepts different from the usual ones will become intelligible to him. (PPF §366)
>
> Then is there something arbitrary about this system? Yes and no. It is akin both to what is arbitrary and to what is non-arbitrary. (Z §358)

The claim that our conceptual scheme is not answerable to reality does not entail that we make arbitrary choices about the concepts we use to represent reality. Indeed, in a great many cases, we do not choose at all. Our forms of representation must be applicable in the circumstances we find ourselves in. They must be applicable to the world by us. This gives us two ways in which it is non-arbitrary: both the conditions of the world (natural laws and the like) and our own

[419] Much of this chapter comes from Couldrick (2022).

cognitive capacities, or, more broadly, our nature, affect our modes of representation. The broader conception of 'our nature' incorporates another aspect of ourselves that affects how we come to represent the world: our goals, purposes, and values.[420] Our nature and the nature of the world we live in can provide us with intelligible reasons for using different concepts. Moreover, they are responsible for those forms of representation that we can hardly be thought of as having chosen to use.

An example might be illustrative of what I have in mind. Imagine two communities that use the term 'gold' for the same kinds of material, found in the same kinds of places. Their criteria for it being 'gold' are primarily based on appearance and feel: 'gold' is yellow, shiny, and reasonably strong. Suppose that one community uses 'gold' to build temples to honour their Gods, whilst the other develops shiny sculptures (the shinier the better) which they bury with their dead as gifts for the deceased to present God with. Now suppose that both communities develop the capacity to atomically analyse their gold. The atomic analysis shows that what they call 'gold' has some subtle variations at the atomic level. To us, this discovery might see us develop further concepts. We might say things like 'what we thought of as one type of material was in fact two', and our rationale for saying this would be that their having (relevantly) different atomic structures entails that they are different materials. Moreover, the analysis may provide us with a new criterion for 'gold', thereby altering the existing concept. Need the two communities share our reaction? Plausibly not. What we do in science may be of little concern to them, or at least of little concern to them in this instance. So, the two communities might shrug their shoulders and say something like 'there is variation in gold just like there is variation in human beings'. Or they might say: 'we have found out that there are, in fact, two kinds of gold!'

But we can equally imagine the communities might be moved to alter their concepts. If the sculptures are presented to God because they are the shiniest objects the community has and, upon separating the materials by atomic structure, they realise one of the two materials known as

[420] In this way, one might detect a kernel of truth when Carnap (1932 [1959]) describes metaphysics as expressing an attitude towards life. But only a kernel, for it is not plausible that all attempts to provide answers to philosophical problems do themselves express attitudes towards life (Hacker 2001, 338). Nor is it plausible that this provides an exhaustive understanding of metaphysics, even for those sympathetic to the idea that metaphysics presented as a true description of reality is incoherent. For metaphysics can also be interpreted as an attempt to propose (sometimes radically) new modes of representation or as a means of clarifying existing ones. Indeed, this understanding of traditional metaphysics is one that receives significant attention in Moore's study of the evolution of modern metaphysics (Moore 2012). As for the kernel of truth in Carnap, given that our favoured modes of representation are driven by the kinds of values and purposes we have, a metaphysical system may well be said to express a possible attitude towards life that one could take (whether it expresses that of its author is another matter). One place I think this a quite natural suggestion is the debate regarding personal identity. One need only reflect on certain Buddhist conceptions, and related conceptions that have been developed in analytic philosophy (e.g. Parfit 1984), to see how one might sympathise with Carnap's suggestion.

'gold' is shinier than the other, then perhaps that community would want to reserve their term 'gold' for the material they offer to God. We could imagine them praying for forgiveness for having sent their loved ones without the appropriate present, without 'gold'. As for the other community, suppose it came to their attention that one of these two substances was substantially stronger than the other and so, out of respect for their God, or concern for the safety of those that practice inside the temples, they choose to only build temples out of the stronger material. This might lead to some conceptual innovation – do they now admit of two types of 'gold', one of which they build with, the other they do something else with? Maybe. Or maybe they will come up with a new title for the weaker 'gold'. Or perhaps they will judge that they need not do anything with the concept but admit of subtle variations in 'gold', where some is better for building than others. Just as some people are better suited to building temples than others. They remain human beings. 'Gold' remains 'gold'. If we were to insist that these were different materials, perhaps they would laugh and tell us we were crazy. 'Look at them side by side', they might say, 'can't you see they are the same thing?'[421]

In this example, the communities seem to have some sort of decision to make. But notice firstly that they might not feel that way. The role 'gold' plays in the life of those that bury the dead with sculptures may make them feel compelled to represent the other substance with a new concept. Of course, they might philosophically reflect and come to realise that there was no strict compulsion to invent a new concept, but they may still rightly say that it is natural for them to do so. It is their way of going on (as it could be for us, albeit for different reasons). If 'gold' is what they gift to God, then the material they no longer present to God is not 'gold'. Secondly, there are some more obvious examples where it is wholly inappropriate to talk of decisions. In a community of people born blind, visual distinctions, i.e. certain modes of representation, are just unavailable. They are inapplicable in the lives of those people. This would be different were some of the community sighted.[422]

The stories just provided illustrate the different ways our concepts come to be conditioned by our own nature and that of the world around us. Our interest in science leads us

[421] LaPorte (2003, 94-100) argues that different cultures' relationships to jade stone provides a real-world example which illustrates much the same point as I have tried to here. See also Donnellan (1983).

[422] I see no reason for thinking those born blind would need to conclude that our utterances lack meaning (after all, those in our community do not draw that conclusion). While our way of speaking is, in some sense, unavailable to them, it does not mean they must conclude we are speaking nonsense – as if the possibilities are bound by their own abilities. They could tell by the way we speak, the way we live with those words and act as a result that they are meaningful expressions. We can have evidence for perceptible differences even if we cannot perceive them ourselves: blind tasters with more finely tuned palates may prove their ability to tell one vintage from another, and the colour-blind may recognise others' capacity to tell green from red via the ability to distinguish ripe and unripe strawberries by sight.

to conceptual structures that pay heed to the way substances interact with one another (which we have learnt is partly determined by their atomic structure). If atomic structure had not played this role, we might have found a different way of characterising those materials. Atomic structure itself might have been conceived of differently, for its status as an underlying nature is in part down to its importance in determining macroscopic properties. The imagined communities of religious people also build conceptual distinctions based on how things are (the appearance or strength of the materials), but an important role was given (explicitly) to their interests and values. While our community's interest in science plays a critical role in how we think about substances, their religious beliefs provided (potential) reasons to alter their ways of thinking. The final example of a community of blind people was used to make explicit the role our physical capacities play in conditioning our conceptual resources. But although I have used examples in which these different aspects have been pronounced, it should be clear that our nature, and the nature of the world around us, condition our sense-making *together* in many, often subtle, ways. What I care to attend to, what I can attend to and what there is to attend to are all sources of reasons (and causes) for the development of the grammar we have. Some may be subject to explicit argument within a community, while others may have more straightforwardly causal links (such as our eyesight) that are seldom grounds for controversy.

2. Prisoners of language?

While the above is widely (though not universally) accepted as part of Wittgenstein's conception of language and the relation it bears to the world, it has also sparked interpretative controversy. The thought that our ways of making sense might be determined by these factors (rather than subject to absolute justification by dint of super-empirical structures) has led some to see the limits of language as limitations. That is, to see us as captives of language, restrained by what Wittgenstein occasionally called our 'form of life' (*Lebensform*).[423] One can usefully compare this (mis)interpretation to how Kant thinks of his transcendental structures imposing necessary constraints on all experience, where creatures like us must represent the world with the same essential aspects. While we are compelled to impose such aspects on our experience (indeed, they make the kind of experience we have possible), they have no bearing on how the world is *in itself*, which is a way of knowing the world that is unavailable to finite creatures like us. In fact,

[423] Hacker (2015) gives an overview of the usage of this phrase in Wittgenstein and the various pitfalls that may arise from exaggerating its role, one of which we shall encounter in this section. A history of the phrase that predates Wittgenstein can be found in work Hacker discusses (e.g., Helmreich and Roosth 2010; Gaffal 2011; Abreu e Silva Neto 2011).

proponents of the 'prisoners of language' interpretation tend to make exactly this comparison (as we shall see below).[424] By claiming that forms of life influence the conceptual resources that are in use, we are said to grant possibilities that we cannot ourselves describe. The thought is that rules for the use of our concepts, such as '5 + 7 = 12' or 'an object cannot be red and green all over simultaneously', hold because of those determining factors (our form of life, which is determined by the kinds of creatures we are and the kind of world we live in).[425] But that implies, so the argument goes, 5 + 7 might not have made 12, had we a different form of life. To acknowledge this possibility, however, would be to lose the necessity contained within those statements. The only option left is to suggest that while there may be alternatives to '5 + 7 = 12', they are not alternatives *for us*. Thrust upon the spearhead of a Kantian contradiction, we must accept that there exist possibilities that we cannot grasp for ourselves, owing to our form of life and the language that grows from it.[426] Our form of life restricts our access to such possibilities, but they do nonetheless exist.

The above enacts what is – in effect – a *reductio ad absurdum* of the position. Hence, we arrive at what I called the 'Kantian contradiction'. We are forced into the absurd position of having to say there are possibilities that we (necessarily) cannot make sense of. We are made to

---

[424] Hanna (2017) makes this comparison most explicitly.

[425] A similar position is entertained by Horwich (2000, 168) when discussing possible sources of *a priori* knowledge, only he focuses on the notion that our minds have an innate structure (which is perhaps more Kantian than the Kantian interpretation of Wittgenstein, though the structure of our minds surely does play some role in determining our form of life). Here's Horwich: 'A further potential source of *a priori* knowledge is the innate structure of our minds. Suppose that each human being is born with, and stuck with, a simple language of thought (i.e. mentalese) containing, amongst other things, certain symbols whose intrinsic nature is such that the principles of classical logic are obeyed. In that case we would have an *a priori* commitment to classical logic. For this commitment would neither derive from experience nor be revisable in light of it.'

[426] Of course, whether Kant is implicated in such a contradiction is debatable. As Sacks (1997, 169) points out, one reading of his transcendental idealism might make him a realist in these terms, through his invocation of transcendental structures and things-in-themselves. One might think it curious, therefore, that some interpreters label the position they attribute to Wittgenstein 'transcendental idealism', for it seems to exclude the doctrine's founder. I would, however, suggest that this curiosity is a greater reflection of a tension within Kant's own thought than it is the result of infelicitous scholarship. For, as Priest (2002, 81) notes, it is not so much *epistemic* access to the noumenal realm that we lack but *conceptual*. Likewise, Cavell (1962 [2002], 65) summarises what he takes to be Wittgenstein's likely response when he writes: 'it would be an illusion not only that we do know things-in-themselves, but equally an illusion that we do not (crudely, because the concept of "knowing something as it really is" is being used without a clear sense, apart from its ordinary language game).' While Kant could not put the objection in these terms, a similar observation might be made from inside his philosophy. Ultimately, Kant concludes that thoughts about things-in-themselves can be true or false (A820–831/B848–859), but whether his readers should agree with this conclusion after accepting that we lack the appropriate concepts to know things-in-themselves (an equally Kantian thought) is an open question. And Kant's task is made considerably more difficult by the fact that he seems to grant (some) knowledge of things-in-themselves, namely the kind of knowledge he expects us to acquire from reading the first *Critique*. Is it possible to acknowledge the truth of Kant's transcendental idealism without doing the very thing that it says is impossible? Adrian Moore (1997; 2012) has devoted much attention to this tension in Kant's thought, and I think its presence warrants the use of 'transcendental idealism' in the context of this chapter so long as we are clear about the differences. Even if the aforementioned contradiction is not Kant's, it remains deeply Kantian. For a response to Moore and arguments against thinking Kant's project is incoherent in this way, see Gardner (2015) (and Moore (2015) for a reply).

argue for the possibility of something while also admitting that 'something' is here entirely empty for us. We are both required to allude to the possibility of representing that which stands outside our powers of representation, *and* – at the same time – expected to maintain that they are not really possibilities (for us) at all. The account thereby draws the limits of language as limitations, as if we were bound from nothing in particular (and yet still, something).[427] The incoherence of such a project resides in its attempt to acknowledge alternatives where it holds that there are none.

It is of critical importance, therefore, that Wittgensteinian accounts can stave off these Kantian (mis)interpretations. That the understanding of the relation between our form of life and conceptual resources is not such that this *reductio ad absurdum* can run its course. It seems Wittgenstein himself recognises this threat when he warns: 'The great difficulty here is not to present the matter as if there were something one *couldn't* do'.[428] Or, as Mulhall, puts it, not to treat the limits of language as if they 'fenced us in, keeping us out of a domain beyond the domain whose limits they stake out, as if nonsense were a peculiar kind of sense, or as if there were something we cannot do here'.[429] The challenge is to avoid conceiving of that which stands outside the bounds of sense as anything more than combinations of words 'excluded from the language, withdrawn from circulation'.[430] Many have rightly pointed out that there are important affinities between Kant and Wittgenstein.[431] The question, then, is how deep do they run?

There has been a trend, inspired by Bernard Williams and Jonathan Lear, to think these run very deep indeed.[432] These interpretations tend to focus on Wittgenstein's use of 'we' throughout his later work. When claims are made regarding how *we* make sense, or what *we* would say, to whom does 'we' refer?[433] Williams, though he recognises its use varies considerably, characterises some instances of 'we' as pointing towards transcendental idealism. The transcendental idealist interpretation sees the possibility of the 'we' expanding to infinity. One might say that it feels the pressure to expand in order to retain the necessity in mathematical statements and the like. The 'we' expanding to infinity makes 'we' include all makers of sense,

---

[427] I am reminded of something Anthony Kenny (2010, 621) says regarding Kant: '[He] is emphatic that it is false to say that there is nothing other than appearance; but to many of his readers it has seemed that a nothing would do just as well as a something about which nothing can be said.'
[428] PI §374.
[429] Mulhall (2009a, 399).
[430] PI §500. Cf. Hacker (2013, 37). This was briefly discussed in chapter one.
[431] Glock (1997), Cavell (2002), Dilman (2002), Hacker (1997, 2013).
[432] Williams (1981); Lear (1982; 1984). Forster (2004), who is broadly sympathetic to Wittgenstein's project and critical of Williams and Lear, also cannot avoid the conclusion that there is something Kantian about his work. Like Hyman (2007), I remain unconvinced.
[433] For my answer to this question, see chapter eight. Here I am primarily concerned with unravelling the mistakes of this interpretation.

thereby forcing statements contrary to necessary truths beyond the bounds of sense, thus protecting the necessity of such truths. This might entail a form of realism: if 'we' were all *possible* makers of sense, then those essential aspects of our representations would just be *the* essential aspects of representation. But there is a second possibility. 'We' might incorporate all *recognisable* makers of sense. For us to be able to represent someone as using language, making sense, they must share our essential aspects of representation. But it wouldn't follow that there were no other modes of representation, only that we could not recognise or represent them as such. So, there are, ultimately, alternatives to our language, but they cannot be represented by us.[434]

One could be forgiven here for thinking this a dead issue. I have already suggested that Wittgenstein thinks forms of representation other than our own can be made intelligible. Indeed, Wittgenstein is famous for including such examples throughout his work.[435] 'Thus', one might submit, 'it is just false that Wittgenstein thought us unable to represent different modes of representation.' In this vein, Mulhall argues that one mistake Williams makes is thinking that 'there is a single, fundamental or determining pattern of use of "we"' in Wittgenstein's work.[436] Moore replies emphatically by quoting Williams who describes the 'pervasive vagueness and indefiniteness evident in the use Wittgenstein makes of "we"'.[437] So where does Mulhall's reading come from? Williams seems to all but reject the only alternative use of 'we' that he explicitly acknowledges, that being the distinction between actual groups of human beings. Williams says the answer to whether Wittgenstein is 'really thinking at all in terms of actual groups of human beings' is 'basically "no"' and that this alternative way of reading 'we' is 'basically misleading'.[438] Williams, then, appears to acknowledge potentially different uses of 'we' only to say that at the very least they are unimportant, and at most they are misinterpretations. But Williams's claims are themselves misleading. For not only will Wittgenstein talk of 'us', but also 'them'. That is,

---

[434] This clearly echoes Stroud (1968). Stroud suggests that although Wittgenstein argued for the possibility of radically different conceptual schemes through examples, these examples are in the end unintelligible. Nevertheless, they succeed in showing us that such radical departures from our own scheme are possible. But how could one succeed in showing the overall possibility of radically different conceptual schemes by providing a set of unintelligible examples? If I make an argument on the back of unintelligible cases, you will rightly think that I have failed to make my case. My case cannot rest on examples that are unintelligible, for those examples are thereby void of any content that could be thought of as supporting anything. I am, therefore, in agreement with Forster (2004), who suggests that Stroud's arguments support Lear's and Williams' interpretation of Wittgenstein's philosophy as deeply (and problematically) Kantian. Coliva (2010) appears to take a similar line to Stroud, and likewise argues that even though we cannot make intelligible radically different practices, our exploration of them does at least show the contingency of our own. But just as I think Forster is right with respect to Stroud, I agree with Kusch (2013, 48) that 'the realization that we cannot imagine radically different practices cannot justify the thought that our practices are contingent.'

[435] Think of examples such as the tribe that uses a concept of pain different from our own (Z, §380), or the community whose measuring sticks expand when heated (RFM I, §5).

[436] Mulhall (2015, 151).

[437] Moore (2015, 348); Williams (1981, 147).

[438] Williams (1981, 159-160).

Wittgenstein employs 'we' in the way Williams claims it is misleading to read him as doing.[439] Moreover, Williams doesn't deny this. Indeed, he explicitly recognises it.[440] Williams describes a 'persistent uncertainty in the interpretation of "we"' introduced by Wittgenstein whenever he talks in terms of actual groups of human beings.[441] Williams thinks these 'imagined alternatives are not alternatives *to* us; they are alternatives *for* us, markers of how far we might go and still remain, within our world'.[442] Williams characterises these alternatives as means by which we experiment within the bounds of our world, within the bounds of sense. 'They' are, ultimately, a part of 'us'. What Williams must find misleading, then, is the idea that 'they' are genuine others. By representing them as such, we thereby disintegrate their otherness.

But where in Wittgenstein is the pressure to expand 'we' to infinity? The central turn in Williams's account comes very late in his essay. He suggests that 'something has to be done if we are to avoid even empirical idealism' for the simple reason that Wittgenstein thinks a human practice or decision is responsible for what we will count as true or false.[443] The distinction between empirical and transcendental idealism goes back to Kant, but it has since been developed by Williams and Moore.[444] According to the understanding relevant to Wittgenstein's philosophy, empirical idealism allows for the alteration of the truth of 'necessary propositions' as a result of human decisions (thereby rendering those truths *unnecessary*). It is, therefore, less so an account of necessity, more an expression of scepticism regarding it.

We need not spend much time on empirical idealism, for a response has already been provided in chapter four. Kalhat objected to the normative account of necessity on the basis that one could not infer from the fact that Kant is a bachelor that he is unmarried, because the relevant norm determining the meaning of 'bachelor' might not have been in force. This was shown to be a misunderstanding. While the sequence of letters might mean something different in another language, the relevant concept 'bachelor' is such that it guarantees that if someone is a bachelor, they are unmarried. While the statement 'a bachelor is an unmarried man' might be counted false in a different language, the proposition expressed by those words in our language cannot be. For it is a norm that helps define the concepts involved.

Empirical idealists replace the realist's logical super-structure with human action. The additional gloss in the context of Wittgenstein is that human action is an aspect of one's form of

[439] PI §200, §207; Z §380, §388; OC §609.
[440] Williams (1981, 155-156).
[441] *Ibid*, 159.
[442] *Ibid*, 160.
[443] Williams (1981, 163).
[444] Williams (1981); Moore (2007; 2012).

life. But Wittgenstein is not minded to replace the logical super-structure with anything. Rules of representation are not in need of support. The standards we employ govern the correct use of our concepts. That we apply those standards is constitutive of our using those concepts. But it does not vindicate them. It merely says the relevant rules are in force and the concepts in use. Likewise, when we stop using them, we do not contest the *truth* of those rules, rather we reject the concepts those rules constitute. Moore makes this clear in his critical notice of Nagel's *The Last Word*.

> Our belief that twice four is eight is indeed necessary. It is necessary because, in 'having this belief', we are effectively laying down a rule of representation: nothing is to count as a disjoint pair of quartets unless it collectively counts as an octet. But this is not in any conflict with the observation that there are all sorts of identifiable contingencies that make it possible for us to have the rule, nor therefore with the observation that we might not have had the rule. If we had not, twice four would not have failed to be eight. Rather, the question of what twice four is would not so much as have arisen for us. We would not have thought in those terms. (Moore 1999 [2019], 138)

What this means is that the contingency Wittgenstein finds at the source of necessity does not entail the empirical idealist's wholesale rejection of it. That rejection is the result of 'necessary truths' being changeable with the actions of language users. Wittgenstein, however, is committed to no such thing. He is keen to emphasise that it is perfectly intelligible to imagine us living under different rules, finding significant different similarities, and having different ways of going on 'in the same way'. What is crucial to distinguishing Wittgenstein's position from idealism is that he recognises that such differences would entail the use of concepts different from our own. In using different concepts, we would be making sense of things differently, as opposed to the empirical idealist's picture where 'necessary truths' change (thus failing to be necessary). Wittgenstein's ability to do this is a result of his insistence that it is the established practice associated with the content of a rule that acts as a standard for correct use, not our continuing use of it. While to become established a rule depends on the actions of language users, once established it does not. It stands independently of those users and actions can be judged against the rules of that practice long after they have given up on it. Rules of representation do not answer to their users (or anything else), though their users may change the rules they follow. But changing the rules one follows does not amount to a denial of the old rules or what followed from them. It is simply a *rejection* of them (a refusal to think in those terms).

Williams's assumption that something must be done to avoid empirical idealism neglects all of this.[445] One reason for the omission is evident in something Williams comes to at the end of his essay. He presupposes there is something to which our language answers, as though a rule requires a guarantee. But I have argued – with Wittgenstein – that no such thing is needed, for the function of a rule is enough insofar as it establishes a constitutive relation between itself and the concepts involved. Now, one might reply: 'of course Williams presupposes that a guarantee is required, for how else is he to attribute (a form of) transcendental idealism to Wittgenstein?' But this rather misses the point. Nowhere are we given a reason for believing that Wittgenstein accepts this presupposition, either in Wittgenstein's work or in Williams's essay. The main passage Williams draws on at this point comes from *Zettel*:

> We have a colour system as we have a number system. Do the systems reside in *our* nature or in the nature of things? How are we to put it? – *Not* in the nature of numbers or colours.[446] (Z §357)

Williams takes from this that Wittgenstein is happy to reject realism but unwilling to reject (or accept) idealism. His unwillingness to accept idealism points in the direction of it being *transcendental,* as it suggests an inability to speak of that which imposes limits on our language, thereby avoiding the acknowledgement of possibilities that lie beyond the bounds of sense, an acknowledgement that, in the end, renders this type of account incoherent. We might understand Williams, then, as saying that Wittgenstein cannot deny that systems reside in our nature, for he thinks it true. Yet he must not affirm it either, for his account is such that it necessarily resists meaningful statement. For what it is worth, independent of what Wittgenstein elsewhere has to say, this reading is plausible. Those of us unwilling to accept the reading must answer two questions: what is the reason for Wittgenstein's hesitation? And why doesn't he reject the idealist's proposal?

One reason is that he rejects the presupposition of the question.[447] Wittgenstein's hesitation speaks to his unease with the terms of the question. If he gives an answer, what does it commit him to? But this doesn't explain why he refrains from rejecting the idealist's proposal

---

[445] These insights may not be open to Williams, for he attributes to Wittgenstein an assertion-conditions theory of meaning. This reading of Wittgenstein has rightly gone out of fashion. Moreover, Williams provides little support for it. Baker and Hacker (1984) refute the reading in their response to Kripke (1982). Cora Diamond (1985) and Meredith Williams (2002) concur with Baker and Hacker. For a dissenting voice: Kusch (2006).

[446] Williams misquotes this passage in his essay, and this might be important for his misreading. Williams has Wittgenstein down as denying that our systems reside in the nature of things. But what is noteworthy is precisely that Wittgenstein does not deny this. As I say below, our practices might be said to reside in the nature of things just as much as in our nature. All Wittgenstein denies here is that our systems reside in the nature of numbers or colours, which if endorsed would be an expression of a kind of Platonism.

[447] A point made by Mulhall (2009a, 396).

having rejected the realist's. Here I think one should consider the ambiguity in the notion that systems reside in our nature. Wittgenstein elsewhere isn't even uncomfortable with the idea that necessary propositions correspond to a reality, so long as you acknowledge that the correspondence 'is of an entirely different kind from what you assume'.[448] And notice that in the passage above he does not deny that our systems reside in the nature of *things,* only that they reside in the nature of numbers or colours. Likewise, we might think there is a sense in which systems have their residence in our nature. Systems must be applicable to the world by us. They reside in our nature, depend upon that nature, insofar as they must be used by us. They are shaped by our cognitive capacities and shared purposes. The systems have their use in our lives and, to some extent, our nature is manifest in the systems we use. They are *our* systems. And they are certainly not independently existing systems. They do not have an essence which stands independent of their established use. Hence Wittgenstein is willing to dismiss the notion that the systems reside in the nature of numbers or colours. None of this need commit him to thinking that our systems *answer to* our nature. But it might explain why he is unwilling to reject the 'idealist's proposal', namely that it doesn't have to be understood as a statement of idealism (any more than the notion of residing in the nature of things need be understood as an endorsement of realism). If we think about the relationship between our form of life (part of which is constituted by our nature) and our concepts, then we might consider the notion that these systems lie, or reside, in our nature as quite apt. It is in our form of life that these systems have their home, outside of it they have no application.[449]

Both Lear and Williams see Wittgenstein's philosophy as an attempt at a kind of self-conscious reflection that leaves us with a greater understanding of the interaction between us and the world. Part of what we are to come to understand is how necessary truths are grounded (made true, accounted for). Or, to put it another way, we are to understand what necessary propositions answer to. So, Lear writes:

> Before we engaged in philosophical reflection, we were disposed to make various assertions, for example, '7+5 must equal 12'. As we study the *Investigations* we come to assert, 'We are so minded as to assert: 7+5 must equal 12'. […] It is such an insight which, I think, led commentators to think that Wittgenstein denied the objectivity of logical or mathematical necessity. However, after we realize that there is (for us) no

[448] LFM 244.

[449] Cf. Z §350. An alternative reading might be that Wittgenstein simply rejects it out of hand on the basis that different languages (such as Russian) employ different colour-concepts from ours. According to this view, it doesn't make sense to say our colour-concepts are determined by our nature because other human cultures (that share our nature) have developed different colour-concepts.

> alternative possibility of being 'other minded' – that is, that there is no alternative possibility – we seem to come back to our original assertion: '7+5 must equal 12'. (Lear 1984, 238).

Lear commits a mistake we should now be familiar with. 'We are so minded as to assert…' makes it sound as though 7 + 5 might not have made 12. But this is not how Wittgenstein sees things. 7 + 5 = 12 of necessity, whether we are so minded as to assert it or not. Lear suggests that our form of life constitutes our being so minded.[450] Thus, he claims our form of life is such that we assert '7 + 5 = 12'. And he is right. But he is right in a much less interesting way than he thinks. We assert '7 + 5 = 12' because it follows from the concepts we use. It is not right *because* we assert it, nor is it right *because* we could not assert anything else.[451] Our inability to assert anything else is a result of it being one of our norms, not a reason for it being one. To assert something else would be to fail to make sense. This isn't to say that our rules do not, in the sense I outlined above, reflect our nature by virtue of their having to be used by us. It is just to say that our nature is not what makes the rule correct or necessary. Nothing does that. The 'rule qua rule is detached; it stands as it were alone in its glory'.[452] Our nature may help determine which rules we use, but it does not give those rules their necessity. As Rush Rhees writes of mathematics, there are '*various* reasons […] for our doing it the way we do, although none of these shows that there is any inherent *necessity* for doing it in this way'.[453]

The mistake the Kantian interpretation makes is to think that such reasons account for, or – in some sense – impose the necessity present in our language. As though the explanation for why we came to use a concept 'aunt' that applies only to females could explain why an aunt *must* be female. But such an explanation is hopeless, not least because the kinds of things those explanations appeal to are unsuitable candidates for guaranteeing necessity. For they are contingent facts that might have been false, thus their use as explanations for necessity reopens the door to the absurdity that aunts might not have been female and 5 + 7 might not have made 12. In fact, once properly understood, the contingency of those explanations merely leaves the door open to the possibility that we might have come to follow different rules, and so we may have lacked a concept 'aunt' where an aunt must be female. As Wittgenstein writes, in a slightly different version of a remark quoted at the beginning of this chapter:

---

[450] Lear (1984, 229).
[451] Cf. RFM VI, §23. It is not the fact that we're trained to accept a given mathematical proposition that matters but what it is we are trained to accept.
[452] RFM VII, §3.
[453] Rhees (1970, 120).

> I am not saying: if such-and-such facts of nature were different people would have different concepts (in the sense of a hypothesis). But: if anyone believes that certain concepts are absolutely the correct ones, and that having different ones would mean not realizing something that we realize – then let him imagine certain very general facts of nature to be different from what we are used to, and the formation of concepts different from the usual ones will become natural to him. *'Natural', not 'necessary'* [my emphasis]. (RPP I, §49)

As we saw in chapter four, we can think of this as there being internal and external aspects of linguistic norms. The internal aspect has the character of necessity: the constitutive relation that obtains between a concept and norm. Its external counterpart, on the other hand, acknowledges the possibility that different norms could be followed, and different concepts used. To use a given concept, one must apply it according to its criteria. But there is no obligation to use that concept and so apply those criteria.

3. The possibility of explanation

There is an important lesson to be taken from Kantian conceptions of Wittgenstein's philosophy that is directly relevant to the question of how the empirical realm relates to the conceptual one. Wittgenstein famously rebuked philosophical explanations when it came to a word's meaning. He was clear that such a task is hopeless, for a word's meaning does not stand in need of explanation but *description*. For what could a philosophical explanation provide? What is left to be uncovered once one understands what one's words mean? It betokens a philosophical confusion to demand to know *why* an aunt must be female. For the answer can stretch no further but to say: because that is just what an aunt *is*, what it means to be an aunt. One wants to say: there is no concept 'aunt' without such a rule. It does not exist independently of the rules that constitute it.[454] This is the point at which Wittgenstein thought explanations come to an end. Our spades turn for the questioner is simply running up against a rule of language. The point is to demystify the 'must', for the 'must' merely masks a rule of the kind: 'this is what we will count as an "aunt"'. We effectively run into a question of identity: the reason one cannot be a male aunt is that in being male one would not be what we call an 'aunt'. If faced with the question 'why?'

[454] As argued in chapter one. Cf. PI, §549.

once more, we could do little better than to say: this is what we mean by 'aunt' and 'male', these are the rules by which we apply those terms.

But the rebuking of philosophical explanations here does not stretch to explanations of all kinds. There may be historical, anthropological, or even scientific reasons for our concepts. I am, therefore, in agreement with Daniel Hutto and Glende Satne that 'it is important to note that Wittgenstein's injunction against philosophical explanation was not a rejection of explanation simpliciter.'[455] I should not like to pre-judge the potential successes other disciplines might have in determining why we came to live, speak, and think how we do, but there is nothing philosophically suspect with thinking that such explanations might exist. Indeed, if one follows the history of science, one will not merely embark on a journey of discoveries, but also a rich history of conceptual innovations made in light of (and occasionally in anticipation of) those discoveries.[456] Many of those innovations have since permeated into everyday discourse, leaving room for both scientific and cultural explanations. The former to explain why those innovations were helpful to science and the latter to explain what it was about our cultural setting and the nature of those innovations that made such permeations occur.[457]

The crucial lesson to be learnt is how these things come apart: whatever reasons there are for our having come to follow the rules that we do, it is the function of a rule, its role in constituting concepts, that accounts for the unmysterious 'must' that has been game to bewilder and inspire awe in equal measure.[458] For those disappointed by this answer, one might ask whether the inspiration that was misguided in respect of the 'must' might be rehabilitated in the acknowledgement of the seemingly infinite sea of possible grammars we could have developed (and might still develop). Whatever we think of the prospects for such a rehabilitation, we should accept that treating necessary truths as immutable laws beyond the boundaries of the language they structure is a mistake and one that, insofar as it limits our imagination and appreciation of possible forms of thought, may not be without practical (and, therefore, ethical) consequence.[459]

---

[455] Hutto and Satne (2018, 157). And, therefore, in disagreement with Malcolm (1993, 82), who suggests neither philosophy nor science can provide an explanation for our form of life.
[456] As touched upon in chapters three and seven. See LaPorte (2003) for a book-length defence of this claim.
[457] Wittgenstein is often touted as being hostile towards science (e.g. in Williams (1981, 159) and Maddy (2014, 122-125)). He is certainly hostile towards scientism and this is evident throughout his philosophy (Beale and Kidd (eds) 2017). His hostility towards science itself, however, is difficult to detect in his philosophy and the impression is perhaps due to 'sporadic cultural reflections' (Mulhall 2009a, 395), often aimed at the social institutions of science rather than the scientific method.
[458] A line Wittgenstein (RFM VII, §67) takes with mathematics is applicable here: 'The mathematical Must is only another expression of the fact that mathematics forms concepts.' The same goes for other inexorable Musts.
[459] There is a burgeoning literature on Wittgensteinian ethics (see, e.g., Crary 2005, De Mesel and Kuusela (eds.) 2021, Diamond 2019, Hacker (2019b), Lovibond 1983, Mulhall 2002, Pleasants 2008). But for something that is especially relevant to the point I make here, consider Hacker's (Baker and Hacker 2014, 330) attempt to make Wittgenstein's wood-sellers intelligible to us (RFM I §149-50; cf. Janik and Toulmin (1973, 228) where another

Moreover, the refusal to acknowledge these possibilities, to steal them away from oneself, betrays a mode of existence that is, it seems to me, essentially characterised by bad faith. For it denies both the responsibility we have for our language and the choices still available to us.

4. Constraint?

Finally, one need not understand conceptual structures as necessarily constraining. One must, after all, have a language in order to speak. And language gives one the tools with which to discuss and criticise a range of things, including the language itself. It provides the logical space in which one can think, but it does not exhaust it. We can develop new ways of speaking. Expressive freedom, following Brandom, is what is given to those who are taught a language, and once acquired it can be put to work in a great variety of ways, including 'in the generation of new possibilities of performance which did not and could not exist outside the framework of norms inherent in the social practices which make up language'.[460] Moreover, there is freedom within a game (or language) as well as in choosing between them, for as Mulhall and Cavell point out, within the rules of a game there is a great range of possibilities to choose from that do not otherwise exist (you cannot score a beautiful goal if there is no such thing as scoring one).[461] There is both a rule book and a coaching manual.[462]

A language may be better or worse than some rival according to a set of criteria, and one may feel constrained by a particular language if it lacks concepts that one deems appropriate or best suited to some purpose. But what I am suggesting is fallacious is the notion that we might do without conditions altogether, as if we might have a language without norms of use that are informed by the contingencies so far discussed.

> [T]he grammar of our words individuates phenomena in ways that express human interests and human nature; the ways in which criteria tell one kind of object from

example in this vein is discussed). Part of our difficulty, Hacker submits, is the example's inconsistency with a crude form of economic motivation. But the logic of the system is unblemished by this fact, and our typical response to it is a result of the scant detail Wittgenstein provides in his sketch coupled with our unwillingness to think seriously about the ways in which other people (possibly our future selves) might live. Some genuine alternatives barely strike us as possible – so we give them not a moment's thought. As Schroeder (2015, 126-7) suggests, it could be that the wood-sellers are stupid, but they might just have different priorities. They might appreciate the skill of spreading logs to cover a large area and be happy to pay more when wood is skilfully arranged, just 'as many of us are happy to pay more for nice packaging or for some affable smiles from the sales staff'. Indeed, Wittgenstein's earlier iterations of the example, where wood is supplied free of charge or the price is always the same no matter the quantity, suggest that he too is challenging us to escape the grip of our picture to see how others might live.

[460] Brandom (1979, 194).

[461] Mulhall (2001, 79) and Cavell (1979, 308).

[462] Cf. McCabe (1994).

> another reflect the distinctions that matter to their users, their shared sense of what is natural and outrageous, what is useful and what pointless. (Mulhall 2001, 260)

The mere existence of rules of representation does not imprison us. Those rules constitute a practice, the mastering of which provides a capacity we otherwise lack. The sense that our language can *require* us to do certain things should not be interpreted as limiting what we can do. By marking where our words have application, norms determine what those words mean – they give them a use. If we indulge Wittgenstein's metaphor of words as tools for a moment, then we shall see that a tool both gives us a means to do something we might otherwise be unable to, and at the same time comes with its own limits (it must be used in a certain way). Crucially, the two are not independent of each another: to do things with words (to take part in the practice made possible by some grammar) is to use those words in the appropriate manner. Of course, having a limited or unsophisticated language can hold us back and the options available to us will be limited by our own capacities. But there is nothing in principle to stop us from adopting new patterns of thought with their own peculiar grammars. Forms of expression enable us to say certain things. They do not stop us from saying something else. Carving wood with a knife does not prevent me from measuring the length of the newly fashioned object. It is my not having something to measure it with that might.

Conclusion

I started this thesis by presenting some problems for substantial accounts of necessary truth, ranging from general epistemological quandaries, to questioning whether statements of them can really be characterised as descriptions of reality. The rest of the thesis built upon the perspective from which those challenges were made and defended an insubstantial account whereby necessary truths are better understood as conceptual propositions. Indeed, chapter two onwards might be seen as an attempt to vindicate the challenges presented to the substantial account in the introduction and first chapter: to show how an insubstantial account can do better than the alternative and make good on some of the claims made at the beginning.

The course the account charts is narrow. First, it must reconcile necessity with its contingent source. I have argued that it can do so by acknowledging internal and external perspectives on a given practice or concept. The internal is characterised by constitutive relations, which are determined by norms. The external focuses on how those norms come about and the fact that we could have inherited and developed a different set. Second, drawing a distinction between norms and constitutive relations on the one hand, and contingent propositions and empirical associations on the other, invites the challenge that this unnaturally constrains language and linguistic possibilities. Yet, so the argument goes, language rarely complies. It is not often easy to distinguish conceptual propositions from their empirical counterparts. However, I have shown how a given proposition may be indeterminate while retaining the distinction in kind between empirical and conceptual relations. Indeed, via the content of our norms, we can explain this indeterminacy while employing the very distinction the examples are supposed to call into doubt. Finally, I have argued that while conceptual propositions do not answer to the world in the way empirical descriptions do, they can nevertheless be justified by reference to reality. But that is by reference to natural (contingent) facts, not some super-empirical structure that we aim to mirror with our conceptual scheme. How the facts intersect with our interests, needs and values can justify the norms of description we employ. Genealogy can bring this out: it can show some of the contingent conditions that may explain how we came to represent the world in the terms we do. That they are contingent shows how things might have been different and may yet be different in the future (a justification that holds now may be superseded later).

I have also suggested that philosophy which understands itself in these terms, as not being concerned with tracing the contours of a hidden reality, need not be considered an entirely

trivial enterprise. In any case, I do not see the point of preserving a myth that is intended to provide philosophy a 'credible' subject matter, other than to invite scorn and ridicule from those that never saw any value in philosophy anyway. As Warren suggests, mistaking practical questions for factual ones hinders progress, insofar as we find ourselves arguing that some given conceptual norm is *really* true rather than evaluating the concepts thereby articulated on pragmatic grounds.[463] But whatever the basis of necessary truth, philosophy gains nothing by pretending conceptual matters are fundamentally trivial. It merely reinforces the prejudices of those who seek to denigrate it. Philosophy, as a minimum, should help us to think clearly, whether that is a matter of distinguishing questions that can be answered empirically from those that can't, clarifying instances of conceptual conflation and confusion, or indeed making certain arguments for representing reality in a particular way (though these arguments are by no means without input from the empirical sciences). If philosophers are not prepared to defend the value of these practices (which can themselves be of value to other disciplines), they should not complain when their subject is inevitably cast as a relic of an unscientific past.

[463] Warren (2022, 52).

References

Works by Wittgenstein

Wittgenstein, L. (1967), *Zettel (Z)*, Anscombe and von Wright (eds.), Anscombe (trans.), Oxford: Basil Blackwell.

Wittgenstein, L. (1969), *The Blue and Brown Books (BB)*, Oxford: Blackwell.

Wittgenstein, L. (1974), *Philosophical Grammar (PG)*, Rhees (ed.), Kenny (trans.), Basil Blackwell.

Wittgenstein, L. (1975), *On Certainty (OC)*, Anscombe and von Wright (eds.), Paul and Anscombe (trans.), Oxford: Basil Blackwell.

Wittgenstein, L. (1976), *Lectures on the Foundations of Mathematics (LFM), Cambridge 1939*, from the notes of Bosanquet, Malcolm, Rhees, and Smythies, Diamond (ed.), Sussex: The Harvester Press.

Wittgenstein, L. (1978), *Remarks on the Foundations of Mathematics (RFM)*, Anscombe, Rhees, and von Wright (eds.), Anscombe (trans.), Oxford: Basil Blackwell.

Wittgenstein, L. (1979), *Remarks on Colour (RC)*, Anscombe (ed.), McAlister and Schättle (trans.), Blackwell.

Wittgenstein, L. (1980), *Remarks on the Philosophy of Psychology (RPP)*, Vol. 1, Oxford: Blackwell.

Wittgenstein, L. (1980), *Wittgenstein's Lectures, Cambridge 1930-32*, *from the notes of J. King and D. Lee (LWL)*, Lee (ed.), Oxford: Blackwell.

Wittgenstein, L. (2005), *The Big Typescript (BT)*, Luckhardt and Aue (eds. and trans.), Oxford: Blackwell.

Wittgenstein, L. (1953 [2009]), *Philosophical Investigations (PI)*, Anscombe, Hacker, and Schulte (trans.), Revised 4th edition by Hacker and Schulte, Oxford: Wiley-Blackwell.

Wittgenstein, L. (2009), *Philosophy of Psychology – A Fragment (PPF)*, Anscombe, Hacker, and Schulte (trans.), Revised 4th edition by Hacker and Schulte, Oxford: Wiley-Blackwell.

Wittgenstein, L. (2016-), *MS 137 (1948-49)*, Interactive Dynamic Presentation (IDP) of Ludwig Wittgenstein's philosophical *Nachlass* [http://wittgensteinonline.no/], edited by the Wittgenstein Archives at the University of Bergen under the direction of Pichler, Bergen: Wittgenstein Archives at the University of Bergen.

Other references

Abreu e Silva Neto, N. (2011), The Uses of "Forms of Life" and the Meanings of Life', in Gálvez and Gaffal (eds.), *Forms of Life and Language Games*, Heusenstamm: Ontos Verlag, 75-106.

Alston, W. P. (1967), 'Vagueness', in Edwards (ed.), *The Encyclopaedia of Philosophy*, Vol. 8, New York: Macmillan, 218–221.

Anscombe, G. E. M. (1981), 'The Question of Linguistic Idealism', *Collected Philosophical Papers*, Vol. 1, Oxford: Basil Blackwell, 112-133.

Appiah K. A. (2008), 'Experimental Philosophy' (Presidential Address to the Eastern Division of the American Philosophical Association 2007), *Proceedings and Addresses of the American Philosophical Association* 82, 7–22.

Austin J. L. (1957). 'A Plea for Excuses', *Proceedings of the Aristotelian Society* 57, 1–30.

Austin, J. L. (1961), 'Other Minds', in *Collected Philosophical Papers*, Oxford: Oxford University Press, 44-84.

Ayer, A. J. (1936), *Language, Truth and Logic*, London: Victor Gollancz Ltd.

Baker, G. (2004), *Wittgenstein's Method: Neglected Aspects*, Morris (ed.), Oxford: Wiley-Blackwell.

Baker, G. and Hacker, P. M. S. (1984), *Scepticism, Rules and Language*, Oxford: Blackwell.

Baker, G. and Hacker P. M. S. (1985), *Rules, Grammar and Necessity*, Oxford: Blackwell.

Baker, G. and Hacker, P. M. S. (1990), 'Malcolm on Language and Rules', *Philosophy* 65, 167-179.

Baker, G. and Hacker, P. M. S. (2014), *Wittgenstein: Rules, Grammar and Necessity*, 2nd edition, extensively revised by Hacker, Oxford: Wiley Blackwell.

Balaguer, M. and Horgan, T. (2016), 'A Solution to the Paradox of Analysis', *Analysis* 76, 3-7.

Bangu, S. (2019), 'Hard and Blind: On Wittgenstein's Genealogical View of Logical Necessity', *Philosophy and Phenomenological Research* 102, 439-458.

Barcan Marcus, R. (1961), 'Modalities and Intensional Languages', *Synthese* 13, 303-322.

Bates, S. and Cohen, T. (1972), 'More on What We Say', *Metaphilosophy* 3, 1–24.

Beale, J. and Kidd, I. J. (eds.) (2017), *Wittgenstein and Scientism*, London: Routledge.

Bird, A. (2000 [2014]), *Thomas Kuhn*, Abingdon: Rouledge.

Bird, A. (2007), 'What is Scientific Progress?', *Nous* 41, 64–89.

Bird, G. (1995), 'Carnap and Quine: Internal and External Questions', *Erkenntnis* 42, 41-64.

Bix, B. (1991), 'H. L. A. Hart and the "Open Texture" of the Language', *Law and Philosophy* 10, 51–72.

Bloor, D. (1996), 'The Question of Linguistic Idealism Revisited', in Sluga and Stern (eds.), *The Cambridge Companion to Wittgenstein*, Cambridge: Cambridge University Press, 354-382.

Boghossian, P. (1996), 'Analyticity Reconsidered', *Nous* 30, 360-391.

Boghossian, P. and Williamson, T. (2020), *Debating the A Priori*, Oxford: Oxford University Press.

Braddon-Mitchell, D. (2008), 'Naturalistic Analysis and the A Priori', in Braddon-Mitchell and Nola (eds.), *Conceptual Analysis and Philosophical Naturalism*, Cambridge, MA: MIT Press, 23-44.

Brandom, R. (1979), 'Freedom and Constraint by Norms', *American Philosophical Quarterly* 16, 187-196.

Brandom, R. (2008), *Between Saying and Doing: Towards an Analytic Pragmatism*, Oxford: Oxford University Press.

Burge, T. (1979), 'Individualism and the Mental', in French, Uehling and Wettstein, *Midwest Studies in Philosophy IV: Studies in Metaphysics*, Minneapolis: University of Minnesota Press, 73-121.

Burgess, A., and Plunkett, D. (2013a), 'Conceptual Ethics I', *Philosophy Compass* 8, 1091–1101.

Burgess, A., and Plunkett, D. (2013b), 'Conceptual Ethics II', *Philosophy Compass* 8, 1102–1110.

Burgess, A., Cappelen, H. and Plunkett, D. (eds.) (2020), *Conceptual Engineering and Conceptual Ethics*, Oxford: Oxford University Press.

Burian, R. (1985), 'On Conceptual Change in Biology: The Case of the Gene', in Depew and Weber, *Evolution at a Crossroads: The New Biology and the New Philosophy of Science*, Cambridge, MA: MIT Press, 29-42.

Büttner, K. M. (2015), 'Is there such a thing as Relative Analyticity?', *Ratio*, 47-56.

Canfield, J. (1981), *Wittgenstein: Language and World*, University of Massachusetts Press.

Canfield, J. (2009), 'Back to the Rough Ground: Wittgenstein and Ordinary Language', in Glock and Hyman (eds.), *Wittgenstein and Analytic Philosophy: Essays for P. M. S. Hacker*, Oxford: Oxford University Press, 109-132.

Cappelen, H. (2018), *Fixing Language: An Essay on Conceptual Engineering*, Oxford: Oxford University Press.

Carnap, R. (1932 [1959]), 'The Elimination of Metaphysics Through Logical Analysis of Language, in Ayer (ed.), *Logical Positivism*, New York: Macmillan, 60-81.

Carnap, R. (1934 [1937]), *The Logical Syntax of Language*, A. Smeaton (trans.), London: Routledge.

Carnap, R. (1947), *Meaning and Necessity: A Study in Semantics and Modal Logic*, Chicago: University of Chicago Press.

Carnap, R. (1950), 'Empiricism, Semantics, and Ontology', *Revue Internationale de Philosophie* 4, 20-40.

Carnap, R. (1955), 'Meaning and Synonymy in Natural Language', *Philosophical Studies* 6, 33-47.

Carnap, R. (1963), 'WV Quine on Logical Truth', in Schilpp (ed.), *The Philosophy of Rudolf Carnap*, 915-922.

Cassam, Q. (1986), 'Necessity and Externality', *Mind* 95, 446-464.

Cavell, S. (1958 [2002]), 'Must We Mean What We Say?', *Must We Mean What We Say?*, Updated Edition, Cambridge: Cambridge University Press, 1-40.

Cavell, S. (1962 [2002]), 'The Availability of Wittgenstein's Philosophy', *Must We Mean What We Say?*, Updated Edition, Cambridge: Cambridge University Press, 41-67.

Cavell, S. (1979), *The Claim of Reason*, Oxford: Oxford University Press.

Cavell, S. (2002), *Must We Mean What We Say*, Cambridge: Cambridge University Press.
Cavell, S. (2010), *Little Did I Know: Excerpts From Memory*, Stanford: Stanford University Press.
Chalmers, D. (2011a), 'Verbal Disputes', *Philosophical Review*, 515-566.
Chalmers, D. (2011b), 'Revisability and Conceptual Change in "Two Dogmas of Empiricism"', *The Journal of Philosophy* 108, 387-415.
Chapman, S. (2011), 'Arne Naess and Empirical Semantics', *Inquiry* 54, 18-30.
Child, W. (2010), 'Wittgenstein's Externalism', *The Later Wittgenstein on Language*, Whiting (ed.), Palgrave Macmillan, 63-80.
Child, W. (2011), *Wittgenstein*, Abingdon: Routledge.
Child, W. (2017), 'Wittgenstein, Scientism, and Anti-scientism in the Philosophy of Mind', in *Wittgenstein and Scientism*, Beale and Kidd (eds.), London: Routledge, 81-100.
Coffa, J. (1991), *The Semantic Tradition: From Kant to Carnap*, Wessels (ed.), Cambridge: Cambridge University Press.
Coliva, A. (2010), 'Was Wittgenstein an Epistemic Relativist?', *Philosophical Investigations* 33, 1-23.
Conant, J. (2001), 'Two Conceptions of Die. Überwindung der Metaphysik', in McCarthy and Stidd (eds.), *Wittgenstein in America*, Oxford: Clarendon Press, 13-61.
Conant, J. (2004), 'Why Worry about the Tractatus?', Stocker (ed.), *Post-Analytic Tractatus*, Aldershot: Ashgate, 167-192.
Cook, J. (1982), 'The Illusion of Aberrant Speakers', *Philosophical Investigations* 5, 215-266.
Couldrick, S. W. A (2022), 'Wittgenstein on Necessity: "Are You Not Really an Idealist in Disguise?"', *Analytic Philosophy* [early view].
Crary, A. and Read, R. (eds.) (2000), *The New Wittgenstein,* London: Routledge.
Crary, A. (2005), 'Wittgenstein and Ethics: A Discussion with Reference to *On Certainty*', in Moyal-Sharrock and Brenner (eds.), *Readings of Wittgenstein's On Certainty*, 275-301.
Dain (2008), 'Wittgenstein, Contextualism, and Nonsense: A reply to Hans-Johann Glock', *Journal of Philosophical Research* 33, 101-125.
Davidson, D. (1973), 'On the Very Idea of a Conceptual Scheme', *Proceedings and Addresses of the American Philosophical Association* 47, 5-20.
De Mesel, B. and Kuusela, O. (eds.) (2021), *Ethics in the Wake of Wittgenstein*, Abingdon: Routledge.
Devitt, M. (1991), *Realism and Truth*, Oxford: Blackwell.
Diamond, C. (1985), 'Review of Baker and Hacker (1984)', *Philosophical Books*, 26-29.
Diamond, C. (1989), 'Rules: Looking in the Right Place', in Phillips and Winch (eds.), *Attention to Particulars*, Palgrave Macmillan, 12-34.

Diamond, C. (1991), *The Realistic Spirit: Wittgenstein, Philosophy, and the Mind*, Cambridge: The MIT Press.

Diamond, C. (2019), *Reading Wittgenstein with Anscombe, Going on to Ethics*, Cambridge, MA: Harvard University Press.

Dickie, I. (2015), *Fixing Reference*, Oxford: Oxford University Press.

Dilman, I. (2002), *Wittgenstein's Copernican Revolution*, London: Palgrave Macmillan.

Donnellan, K. (1962), 'Necessity and Criteria', *The Journal of Philosophy* 59, 647-658.

Donnellan, K. (1983), 'Kripke and Putnam on Natural Kind Terms', in Ginet and Shoemaker (eds.), *Knowledge and Mind*, Oxford: Oxford University Press, 84-104.

Dummett, M. (1959), 'Wittgenstein's Philosophy of Mathematics', *The Philosophical Review* 68, 324-348.

Dupré, J. (1981), 'Natural Kinds and Biological Taxa', *The Philosophical Review* 90, 66-90.

Dupré, J. (1993), *The Disorder of Things*, Cambridge MA: Harvard University Press.

Dutilh Novaes, C. (2020), 'Carnap Meets Foucault: Conceptual Engineering and Genealogical Investigations', *Inquiry*, 1-27.

Edmonds, D. (2020), *The Murder of Professor Schlick: The Rise and Fall of the Vienna Circle*, Oxford: Princeton University Press.

Eklund, M. (2015), 'Intuitions, Conceptual Engineering, and Conceptual Fixed Points', in Daly (ed.), *The Palgrave Handbook of Philosophical Methods*, Basingstoke: Palgrave-Macmillan, 363–385.

Enoch, D. (2007), 'Epistemicism and Nihilism about Vagueness: What's the Difference?' *Philosophical Studies* 133, 285–311.

Ewing, A. C. (1940), 'The Linguistic Theory of A Priori Propositions', *Proceedings of the Aristotelian Society* 15: 207-244.

Faulkner, N. (2010), 'Wittgenstein's Philosophical Grammar: A Neglected Discussion of Vagueness', *Philosophical Investigations* 33, 159-183.

Feyerabend, P. (1988), *Against Method*, Revised edition, London: Verso.

Flanagan, B. (2013), 'Analyticity and the Deviant Logician: Williamson's Argument from Disagreement', *Acta Analytica*, 345-352.

Floyd, J. (2000), 'Wittgenstein, Mathematics and Philosophy', in Crary and Read (eds.), *The New Wittgenstein*, London: Routledge, 232–261.

Fodor, J. and Katz, J. (1963), 'The Availability of What We Say', *The Philosophical Review* 72, 55–71.

Forster, M. (2004). *Wittgenstein on the Arbitrariness of Grammar*, Oxford: Princeton University Press.

Frege, G. (1884 [1968]), *The Foundations of Arithmetic*, J. L. Austin (trans.), Oxford: Basil Blackwell.

Frege, G. (1893/1903 [2013]), *Gottlob Frege: Basic Laws of Arithmetic*, P. Ebert and M. Rossberg (eds. and trans.), Oxford: Oxford University Press.
Friedman, M. (2000), 'Transcendental Philosophy and A Priori Knowledge: A Neo-Kantian Perspective', in Boghossian and Peacocke (eds.), *New Essays on the A Priori*, Oxford: Oxford University Press, 367-383.
Friedman, M. (2001), *Dynamics of Reason*, Stanford: CSLI Publications.
Gaffal, M. (2011), 'Forms of Life as Social Techniques', in Gálvez and Gaffal (eds.), *Forms of Life and Language Games*, Heusenstamm: Ontos Verlag, 57-74.
Gardner, S. (2015), 'Transcendental Idealism at the Limit', *Philosophical Topics* 43, 63-85.
Gettier, E. (1963), 'Is Knowledge Justified True Belief?', *Analysis*, 121-123.
Giaquinto, M. (2008), 'The Linguistic View of a Priori Knowledge', *Philosophy* 83, 89–111.
Glock, H. J. (1997), 'Kant and Wittgenstein: Philosophy, Necessity, and Representation', *International Journal of Philosophical Studies* 5, 285-305.
Glock, H. J. (2002), 'Does Ontology Exist?', *Philosophy* 77, 235-260.
Glock, H. J. (2003a), *Quine and Davidson on Language, Thought and Reality*, Cambridge: Cambridge University Press.
Glock, H. J. (2003b), 'The Linguistic Doctrine Revisited', *Grazer Philosophische Studien* 66, 143-170.
Glock, H. J. (2007), 'Relativism, Commensurability and Translatability', *Ratio* 20, 377-402.
Glock, H. J. (2008), 'Necessity and Language: In Defence of Conventionalism', *Philosophical Investigations* 31, 24-47.
Glock, H. J. (2009), 'Concepts, Conceptual Schemes and Grammar', *Philosophia*, 653-668.
Glock, H. J. (2010), 'From Armchair to Reality?', *Ratio*, 339-348.
Glock, H. J. (2012a), 'Strawson's Descriptive Metaphysics', in Haaparanta and Koskinnen (eds.), *Categories of Being: Essays on Metaphysics and Logic*, New York: Oxford University Press, 391-419.
Glock, H. J. (2012b), 'What is a Theory of Meaning?', *Cahiers Ferdinand de Saussure* 65, 51-79.
Glock, H. J. (2015), 'Nonsense Made Intelligible', *Erkenntnis* 80, 111-136.
Glock, H. J. (2017), 'Impure Conceptual Analysis', in D'Oro and Overgaard (eds.), *The Cambridge Companion to Philosophical Methodology*, 77-100.
Glock, H. J. and Preston, J. (1995), 'Externalism and First Person Authority', *The Monist*, 515-533.
Godman, M., Mallozzi, A., and Papineau, D. (2020), 'Essential Properties are Super-Explanatory: Taming Metaphysical Modality', *Journal of the American Philosophical Association* 3, 1-19.
Grantham, T. (2004), 'Conceptualizing the (Dis)unity of Science', *Philosophy of Science*, 133-155.
Grice, P. (1989), *Studies in the Way of Words*, Cambridge, MA: Harvard University Press.

Grice, P. and Strawson, P. F. (1956), 'In Defense of Dogma', *Philosophical Review* 65, 141-158.
Guess, R. (1994), 'Nietzsche and Genealogy', *European Journal of Philosophy* 2, 274–92.
Gustafsson, M. (2005), 'Perfect Pitch and Austinian Examples: Cavell, McDowell, Wittgenstein, and the Philosophical Significance of Ordinary Language', *Inquiry* 48, 356-389.
Hacker, P. M. S. (1996a), *Wittgenstein's Place in 20th Century Analytic Philosophy*, Oxford: Blackwell.
Hacker, P. M. S. (1996b), 'On Davidson's Idea of a Conceptual Scheme', *The Philosophical Quarterly* 46, 289-307.
Hacker, P. M. S. (1997), *Insight and Illusion: Themes in the Philosophy of Wittgenstein*, Revised edition, Bristol: Thoemmes Press.
Hacker, P. M. S. (2001), *Wittgenstein: Connections and Controversies*, Oxford: Oxford University Press.
Hacker, P. M. S. (2003), 'Wittgenstein, Carnap and the New Wittgensteinians', *Philosophical Quarterly* 53, 1–23.
Hacker, P. M. S. (2004), 'Substance: Things and Stuff', *Supplement to the Proceedings of the Aristotelian Society*, 41-63.
Hacker, P. M. S. (2006), 'Soames' History of Analytic Philosophy', *The Philosophical Quarterly* 56, 121–131.
Hacker, P. M. S. (2007), 'Analytic Philosophy: Beyond the Linguistic Turn and Back Again', in Beaney (ed.), *The Analytic Turn: Analysis in Early Analytic Philosophy and Phenomenology*, London: Routledge, 125-41.
Hacker, P. M. S. (2009), '"A Philosopher of Philosophy", critical notice of Timothy Williamson's *The Philosophy of Philosophy*', *The Philosophical Quarterly* 59, 337-48.
Hacker, P. M. S. (2013), *Wittgenstein: Comparisons and Context*, Oxford: Oxford University Press.
Hacker, P. M. S. (2015), 'Forms of Life, *Nordic Wittgenstein Review*, 1-20.
Hacker, P. M. S. (2019a), *Wittgenstein: Meaning and Mind* (Volume 3 of an Analytical Commentary on the Philosophical Investigations), Part 1: Essays, 2nd Edition, Oxford: Wiley Blackwell.
Hacker, P. M. S. (2019b), *The Moral Powers*, Oxford: Wiley-Blackwell.
Hacking, I. (2007), 'Putnam's Theory of Natural Kinds and their Names is not the Same as Kripke's', *Principia: An International Journal of Epistemology*, 1-24.
Hale, B. (2013), *Necessary Beings: An Essay on Ontology, Modality, and the Relations Between Them*, Oxford: Oxford University Press.
Hanfling, O. (1984), 'Scientific Realism and Ordinary Usage', *Philosophical Investigations*, 187-205.
Hanfling, O. (2000), *Philosophy and Ordinary Language: The Bent and Genius of Our Tongue*, London: Routledge.

Hanfling, O. (2001), 'What is Wrong With Sorites Arguments?', *Analysis* 61, 29-35.
Hanna, R. (2017). 'Wittgenstein and Kantianism', in Glock and Hyman (eds.), *A Companion to* Wittgenstein, Oxford: Wiley-Blackwell, 682-698.
Hansen, N. (2014), 'Contemporary Ordinary Language Philosophy', *Philosophy Compass* 9, 556–569.
Hansen, N. (2017), 'Must We Measure What We Mean?', *Inquiry* 60, 785-815.
Harman, G. (1967), 'Quine on Meaning and Existence I and II', *Review of Metaphysics* 21, 124–151; 343–367.
Harman, G. (2018), *Object-Oriented Ontology: A New Theory of Everything*, Pelican.
Hart, H. L. A. (1961), *The Concept of Law*, Oxford: Clarendon Press.
Haslanger, S. (2012), *Resisting Reality*, Oxford: Oxford University Press.
Helmreich, S. and Roosth, S. (2010), 'Life Forms: a Keyword Entry', *Representations* 112, 27-53.
Hookway, C (1988), *Quine: Language, Experience and Reality*, Cambridge: Polity Press.
Horwich, P. (2000), 'Stipulation, Meaning, and Apriority', in Boghossian and Peacocke (eds.), *New Essays on the A Priori*, Oxford: Oxford University Press, 150-169.
Hutto, D. D. and Satne, G. (2018), 'Wittgenstein's Inspiring View of Nature: On Connecting Philosophy and Science Aright', *Philosophical Investigations* 41, 141-160.
Hyde, D. (2008), *Vagueness, Logic and Ontology*, Aldershot: Ashgate.
Hyde, D. and Raffman, D. (2018), 'Sorites Paradox', *The Stanford Encyclopedia of Philosophy* (Summer 2018 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/sum2018/entries/sorites-paradox/>.
Hylton, P. (2019), 'Carnap and Quine on Analyticity: The Nature of the Disagreement', *Nous*, 1-18.
Hyman, J. (2007), 'Review of Forster (2004)', *The Philosophical Review* 116, 471-473.
Jackman, H. (2001), 'Ordinary Language, Conventionalism, and a Priori Knowledge', *Dialectica* 55, 315–325.
Jackman, H. (2005), 'Intuitions and Semantic Theory', *Metaphilosophy* 36, 363–380.
Jackson, B. B. (2014), 'Verbal Disputes and Substantiveness', *Erkenntnis* 79, 31-54.
Jackson, F. (1998), *From Metaphysics to Ethics: A Defence of Conceptual Analysis*, Oxford: Oxford University Press.
Janik, A. and Toulmin, S. (1973), *Wittgenstein's Vienna*, New York: Simon and Schuster.
Kail, P. (2011), '"Genealogy" and the Genealogy', in May (ed.), *Nietzsche's Genealogy of Morality: A Critical Guide*, Cambridge: Cambridge University Press, 214–33.

Kail, P. (2021), 'Review of Queloz (2021): *The Practical Origins of Ideas: Genealogy as Conceptual Reverse-Engineering'*, *Notre Dame Philosophical Reviews*.

Kalhat, J. (2008a), 'Has the Later Wittgenstein Accounted for Necessity?', *Philosophical Investigations* 31, 1-23.

Kalhat, J. (2008b), 'Necessity and Language: The Gap is Still Very Real', *Philosophical Investigations* 31, 227-236.

Kant, I. (1781 [1998]), *The Critique of Pure Reason*, Guyer and Wood (trans.), Cambridge: Cambridge University Press.

Kaplan, D. (1989), 'Demonstratives', in Almog, Perry and Wettstein, *Themes from Kaplan*, Oxford: Oxford University Press, 481-563.

Kauppinen, A. (2007), 'The Rise and Fall of Experimental Philosophy', *Philosophical Explorations* 10, 95-118.

Keefe, R. (2000), *Theories of Vagueness*, Cambridge: Cambridge University Press.

Keefe, R. and Smith, P. (eds.) (1996), *Vagueness: A Reader*, Cambridge, MA: The MIT Press.

Kenny, A. (2010), *A New History of Western Philosophy*, Oxford: Oxford University Press.

Khatchadourian, H. (2007), *Meaning and Criteria*, New York: Peter Lang Publishing.

Klagge, J. (2017), 'Wittgenstein, Science, and the Evolution of Concepts, in Beale and Kidd (eds.), *Wittgenstein and Scientism*, London: Routledge.

Kneale, W. C. (1947), 'Are Necessary Truths True by Convention?' *Proceedings of the Aristotelian Society*, Supplementary Volume 21, 118–133.

Knobe, J. (2007), 'Experimental Philosophy and Philosophical Significance', *Philosophical Explorations* 10, 119-121.

Knobe, J. and Nichols, S. (eds.) (2008), *Experimental Philosophy*, Oxford: Oxford University Press.

Knobe, J., (2003), 'Intentional Action and Side-Effects in Ordinary Language', *Analysis* 63, 190-193.

Kornblith, H. (2002), *Knowledge and its Place in Nature*, Oxford: Oxford University Press.

Kripke, S. (1980), *Naming and Necessity*, Oxford: Wiley-Blackwell.

Kripke, S. (1982), *Wittgenstein on Rules and Private Language*, Cambridge, MA: Harvard University Press.

Kroon, F. (1987), 'Causal Descriptivism', *Australasian Journal of Philosophy* 65, 1-17.

Kuhn, T. (1962), *The Structure of Scientific Revolutions*, Chicago: University of Chicago Press.

Kuhn, T. (1970 [2000]), 'Reflection on My Critics', in Conant and Haugeland (eds.), *The Road Since Structure*, Chicago: The University of Chicago Press, 123-175.

Kuhn, T. (1976 [2000]), 'Theory Change as Structure Change: Comments on the Sneed Formalism', in Conant and Haugeland (eds.), *The Road Since Structure*, Chicago: The University of Chicago Press, 176-195.
Kuhn, T. (1982 [2000]), 'Commensurability, Comparability, Communicability', in Conant and Haugeland (eds.), *The Road Since Structure*, Chicago: The University of Chicago Press, 33-57.
Kuhn, T. (1983 [2000]), 'Rationality and Theory Choice', in Conant and Haugeland (eds.), *The Road Since Structure*, Chicago: The University of Chicago Press, 208-215.
Kuhn, T. (1990). 'Dubbing and Redubbing: The Vulnerability of Rigid Designation', in Savage (ed.), *Scientific Theories*, Minnesota Studies in the Philosophy of Science, Vol. 14, Minneapolis: University of Minnesota Press, 298–318.
Kuhn, T. (1993), 'Afterword', in Horwich (ed.), *World Changes: Thomas Kuhn and the Nature of Science*, Cambridge, MA: MIT Press, 311-341.
Kusch, M. (2006), *A Sceptical Guide to Meaning and Rules*, Montreal: McGill-Queen's University Press.
Kusch, M. (2013), 'Annalisa Coliva on Wittgenstein and Epistemic Relativism, *Philosophia* 41, 37–49.
Kuusela, O. (2015), 'Wittgenstein and Logic Today: The Logical Must by Penelope Maddy', *Nordic Wittgenstein Review*, 233-236.
LaPorte, J. (2003), *Natural Kinds and Conceptual Change* (Cambridge Studies in Philosophy and Biology), Cambridge: Cambridge University Press.
Lassiter, D. (2008). 'Semantic Externalism, Language Variation, and Sociolinguistic Accommodation', *Mind and Language* 23, 607–633.
Lavers, G. (2019), 'Waismann: From Wittgenstein's Tafelrunde to His Writings on Analyticity', in Makovec and Shapiro (eds.), *Friedrich Waismann: The Open Texture of Analytic Philosophy*, 131-158.
Lear, J. (1982), 'Leaving the World Alone' *The Journal of Philosophy* 79, 382-403.
Lear, J. (1984), 'The Disappearing 'We', *Proceedings of the Aristotelean Society*, Supplementary Volume 58, 219-242.
Lewis, D. (1969), *Convention: A Philosophical Study*, Cambridge, MA: Harvard University Press.
Lewy, C. (1976), *Meaning and Modality*, Cambridge: Cambridge University Press.
Livingston, D. (2011), Review of W. Child's Wittgenstein, *Notre Dame Philosophical Reviews*.
Lorenzini, D. (2020), 'On Possibilising Genealogy', *Inquiry*, 1-22.
Lovibond, S. (1983), *Realism and Imagination in Ethics*, Oxford: Blackwell.
Lowe. E. J. (2007), 'A Problem for A Posteriori Essentialism Concerning Natural Kinds', *Analysis*, 286-292.

Ludlow, P. (2008), 'Cheap Contextualism', *Philosophical Issues* 18, 104-129.
Ludlow, P. (2014), *Living Words: Meaning Underdetermination and the Dynamic Lexicon*, Oxford: Oxford University Press.
Ludwig, K. and Ray, G. (2002), 'Vagueness and the Sorites Paradox', in Tomberlin (ed.), *Language and Mind Philosophical Perspectives* 16, Atascadero, CA: Ridgeview Press, 419–46.
Lugg, A. (2012), 'W.V. Quine on Analyticity: "Two Dogmas of Empiricism" in Context', *Dialogue* 51, 231-246.
Lycan, W. (1994), *Modality and Meaning*, Dordrecht: Springer.
Maddy, P. (2000), 'Naturalism and the *A Priori*', in Boghossian and Peacocke (eds.), *New Essays on the A Priori*, Oxford: Oxford University Press, 92-116.
Maddy, P. (2014), *The Logical Must: Wittgenstein on Logic*, New York: Oxford University Press.
Malcolm, N. (1989), 'Wittgenstein On Language and Rules', *Philosophy* 64, 5-28.
Malcolm, N. (1993), *Wittgenstein: A Religious Point of View?*, Winch (ed.), London: Routledge.
Mates, B. (1958), 'On the Verification of Statements about Ordinary Language', *Inquiry* 1, 161–171.
McCabe, H. (1994), 'Manuals and Rule Books', in Wilkins (ed.), *Considering "Veritatis Splendor"*, Cleveland: Pilgrim Press, 61-68.
McGinn, M. (2013), *Routledge Philosophy Guidebook to Wittgenstein and the Philosophical Investigations*, London: Routledge.
Millikan, R. (1983), *Language, Thought, and Other Biological Categories*, Cambridge, MA, MIT Press.
Moore, A. W. (1997), *Points of View*, Oxford: Oxford University Press.
Moore, A. W. (1997 [2019]), 'The Underdetermination/Indeterminacy Distinction and the Analytic/Synthetic Distinction', *Language, World, and Limits: Essays in the Philosophy of Language and Metaphysics*, Oxford: Oxford University Press, 17-38.
Moore, A. W. (1999 [2019]), 'One or Two Dogmas of Objectivism', *Language, World, and Limits: Essays in the Philosophy of Language and Metaphysics*, Oxford: Oxford University Press, 132-142.
Moore, A. W. (2002 [2019]), 'What are These Familiar Words Doing Here?', in *Language, World, and Limits: Essays in the Philosophy of Language and Metaphysics*, Oxford: Oxford University Press, 39-55.
Moore, A. W. (2003), 'Ineffability and Nonsense', *Proceedings of the Aristotelean Society*, 169-193.
Moore, A. W. (2007), 'Wittgenstein and Transcendental Idealism', in Kahane, Kanterian, and Kuusela (eds), *Wittgenstein and His Interpreters: Essays in Memory of Gordon Baker*, Oxford: Blackwell, 174-199.

Moore, A. W. (2012), *The Evolution of Modern Metaphysics: Making Sense of Things*, Cambridge: Cambridge University Press.
Moore, A.W. (2013), 'Was the Author of the Tractatus a Transcendental Idealist?', eds. Sullivan and Potter, *Wittgenstein's Tractatus: History and Interpretation*, Oxford University Press, 239-255.
Moore, A. W. (2015), 'Replies', *Philosophical Topics* 43, 329-383.
Morris, K. (2019), '"How I See Philosophy": An Apple of Discord Among Wittgenstein Scholars', in Makovec and Shapiro (eds.), *Friedrich Waismann: The Open Texture of Analytic Philosophy*, 107-130.
Moyal-Sharrock, D. (2017), 'Fighting Relativism: Wittgenstein and Kuhn', in Neges, Mitterer, Kletzl and Kanzian (eds.), *Realism - Relativism - Constructivism: Proceedings of the 38th International Wittgenstein Symposium in Kirchberg*, De Gruyter, 215-232.
Moyn, S. (2010), *The Last Utopia: Human Rights in History*, Cambridge: Harvard University Press.
Moyn, S. (2018), *Not Enough: Human Rights in an Unequal World*, Cambridge: Harvard University Press.
Mulhall, S. (2001), *Inheritance and Originality: Wittgenstein, Heidegger, Kierkegaard*, Oxford: Oxford University Press.
Mulhall, S. (2002), 'Ethics in the Light Of Wittgenstein', *Philosophical Papers* 31, 293-321.
Mulhall, S. (2007), *Wittgenstein's Private Language: Grammar, Nonsense, and Imagination in Philosophical Investigations*, Oxford: Oxford University Press.
Mulhall, S. (2009a), 'Hopelessly Strange': Bernard Williams's Portrait of Wittgenstein as Transcendental Idealist, *European Journal of Philosophy* 17, 386-404.
Mulhall, S. (2009b), 'Language-Games and Language: Rules, Normality Conditions and Conversation', in Glock and Hyman (eds.), *Wittgenstein and Analytic Philosophy: Essays for P.M.S Hacker*, Oxford: Oxford University Press, 152-174.
Mulhall, S. (2015), 'Adrian Moore's Wittgenstein', *Philosophical Topics* 43, 149-160.
Mulhall, S. (2018), 'How Complex is a Lemon?', *London Review of Books*, Vol. 40 No.18.
Nado, J. (2019), 'Conceptual Engineering, Truth, and Efficacy', *Synthese* 198, 1507-1527.
Naess, A, (1957) 'Synonymity as Revealed by Intuition', *The Philosophical Review* 66, 87–93.
Naess, A. (1938), *"Truth" as Conceived of by Those Who Are Not Professional Philosophers*, Oslo: I kommisjon hos J. Dybwad.
Nagel, T. (1986), *The View From Nowhere*, Oxford: Oxford University Press.
Needham, P. (2011), 'Microessentialism: What is the Argument?', *Nous*, 1-21.
Nietzsche, F. (1887 [2006]), *On the Genealogy of Morals*, Diethe (trans.), Cambridge: Cambridge University Press.

Nozick, R. (2001), *Invariances: The Structure of the Objective World*, Cambridge, MA: Harvard University Press.

Nyseth, F. (2017), 'Semantic Facts and a Priori Knowledge', *Philosophy* 92, 297-304.

Nyseth, F. (2018), *Language, Necessity and Convention* (doctoral dissertation), retrieved from Apollo – University of Cambridge Repository.

Nyseth, F. (2021), 'Linguistic Conventionalism and the Truth-Contrast Thesis', *The Philosophical Quarterly*, 264–285.

Pap, A. (1958), *Semantics and Necessary Truth: An Inquiry into the Foundations of Analytic Philosophy*, New Haven: Yale University Press.

Papineau, D. (1996), 'Theory-Dependent Terms', *Philosophy of Science* 63, 1-20.

Papineau, D. (2009), 'The Poverty of Analysis', *Proceedings of the Aristotelian Society, Supplementary Volumes*, 1-30.

Parfit, D. (1984), *Reasons and Persons*, Oxford: Oxford University Press.

Peacocke, C. (1999), *Being Known*, Oxford: Oxford University Press.

Peacocke, C. (2005), 'The A Priori', in Jackson and Smith (eds.), *The Oxford Handbook of Contemporary Philosophy*, Oxford: Oxford University Press, 739-763.

Pettit, P. (2020), 'Analyzing Concepts and Allocating Referents', in Burgess, Cappelen and Plunkett, (eds.), *Conceptual Engineering and Conceptual Ethics*, Oxford: Oxford University Press, 333-357.

Pinder, M. (2020), 'On Strawson's Critique of Explication as a Method in Philosophy', *Synthese* 197, 955-981.

Plato (1976), *Meno*, Grube (trans.), Indianapolis: Hackett.

Pleasants, N. (2008), 'Wittgenstein, Ethics and Basic Moral Certainty', *Inquiry* 51, 241-267.

Plunkett, D. (2015), 'Which Concepts Should We Use? Metalinguistic Negotiations and the Methodology of Philosophy', *Inquiry* 58, 828–874.

Plunkett, D. and Sundell, T. (2013), 'Disagreement and the Semantics of Normative and Evaluative Terms', *Philosophers' Imprint* 13, 1-37.

Pole, D. (1958), *The Later Philosophy of Wittgenstein: A Short Introduction*, London: Athlone Press.

Poscher, R. (2011), 'Ambiguity and Vagueness in Legal Interpretation', in Solan and Tiersma (eds.), *The Oxford Handbook of Language and Law*, Oxford: Oxford University Press, 128–144.

Preston, J. (2004), 'Bird, Kuhn, and Positivism', *Studies in History and Philosophy of Science Part A* 35, 327-335.

Preston, J. and Schroeder, S. (2020), 'Explication, Description and Enlightenment', *Logical Analysis and History of Philosophy* 22, Special issue: Analysis and Explication in 20th Century Philosophy, 106-120.

Price, H. (1997), 'Carnap, Quine and the Fate of Metaphysics', *Electronic Journal of Analytic Philosophy* 5.

Price, H. (2009), 'Metaphysics After Carnap: The Ghost Who Walks?', in Chalmers, Manley and Wasserman (eds.), *Metametaphysics: New Essays on the Foundations of Ontology*, Oxford University Press, 320-46.

Priest, G. (2002), *Beyond the Limits of Thought*, Oxford: Oxford University Press.

Putnam, H. (1962 [1975]), 'Dreaming and "Depth Grammar"', in *Mind, Language and Reality: Philosophical Papers* (Vol. 2), Cambridge, MA: Cambridge University Press, 304–324.

Putnam, H. (1962a), 'The Analytic and Synthetic', in Feigel and Maxwell (eds.), *Scientific Explanation, Space, and Time* (Minnesota Studies in the Philosophy of Science, Vol. 3), Minneapolis: University of Minnesota Press, 358-397.

Putnam, H. (1962b), 'It Ain't Necessarily So', *Journal of Philosophy* 59, 658-671.

Putnam, H. (1973), 'Meaning and Reference', *Journal of Philosophy* 70, 699-711.

Putnam, H. (1975), 'The Meaning of "Meaning"', *Mind, Language, and Reality*, Cambridge: Cambridge University Press, 215-271.

Putnam, H. (1983), 'Possibility and Necessity', in *Realism and Reason: Philosophical Papers*, Vol. 3, Cambridge: Cambridge University Press, 46-68.

Putnam, H. (1988), *Representation and Reality*, Cambridge, MA: MIT Press.

Putnam, H. (1990), 'Is Water Necessarily H_2O?', in *Realism With a Human Face*, Conant(ed.), Cambridge, MA: Harvard University Press, 54-79.

Queloz, M. (2021), *The Practical Origins of Ideas: Genealogy as Conceptual Reverse-Engineering*, Oxford: Oxford University Press.

Quine, W. V. O. (1936 [1966]), 'Truth by Convention', *The Ways of Paradox and Other Essays*, New York: Random House, 70-99.

Quine, W. V. O. (1948 [1980]), 'On What There Is', *From a Logical Point of View*, Cambridge, MA: Harvard University Press, 1–19.

Quine, W. V. O. (1951), 'Two Dogmas of Empiricism', *Philosophical Review* 60, 20–43.

Quine, W. V. O. (1960 [2013]), *Word and Object*, New edition, Cambridge, MA: MIT Press.

Quine, W. V. O. (1960a), 'Carnap and Logical Truth, *Synthese* 12, 350–374.

Quine, W. V. O. (1960b), *Word and Object*, Cambridge, MA: MIT Press.

Quine, W. V. O. (1968), 'Ontological Relativity', *The Journal of Philosophy* 65, 185-212.

Quine, W. V. O. (1969), 'Replies', in Davidson and Hintikka (eds.), *Words and Objections*, Dordrecht: Reidel.

Quine, W. V. O. (1974), *The Roots of Reference*, LaSalle: Open Court.

Quine, W. V. O. (1976), *Ways of Paradox and Other Essays*, Revised 2nd edition, Cambridge, MA, Harvard University Press.

Quine, W. V. O. (1986), 'Replies to Critics', in Hahn and Schilpp (eds.), *The Philosophy of W. V. Quine*.

Quinton, A. M. (1954), 'On Punishment', *Analysis* 14, 133-142.

Raatikainen, P. (2020), 'Theories of Reference: What Was the Question?', in Bianchi (ed.), *Language and Reality From a Naturalistic Perspective: Themes From Michael Devitt*, Springer, 69–103.

Railton, P. (2000), 'A Priori Rules: Wittgenstein on the Normativity of Logic', in Boghossian and Peacocke (eds.), *New Essays on the A Priori*, Oxford: Oxford University Press, 170-196.

Rawls, J. (1999), *A Theory of Justice*, Oxford: Oxford University Press.

Ray, G. (2004), 'Williamson's Master Argument on Vagueness', *Synthese* 138, 175-206.

Reichenbach, H. (1920), *Relativitätstheorie und Erkenntnis Apriori*, Berlin: Springer.

Rhees, R. (1970), *Discussions of Wittgenstein*, Bristol: Thoemmes Press.

Rundle, B. (1990), *Wittgenstein and Contemporary Philosophy of Language*, Oxford: Blackwell.

Russell, B. (1905), 'On Denoting', *Mind* 14, 479–493.

Russell, B. (1914 [1917]), 'On Scientific Method in Philosophy', *Mysticism and Logic*, London: Allen and Unwin, 97-124.

Russell, B. (1923), 'Vagueness', *Australasian Journal of Philosophy* 1, 84–92.

Russell, G. (2008), *Truth in Virtue of Meaning: A Defence of the Analytic/Synthetic Distinction*, Oxford: Oxford University Press.

Russell, G. (2019), 'Breaking the Spell: Waismann's Papers on the Analytic/Synthetic Distinction', in Makovec and Shapiro (eds.), *Friedrich Waismann: The Open Texture of Analytic Philosophy*, 159-187.

Ryle, G. (1949), *The Concept of Mind*, London: Routledge.

Ryle, G. (1950), "'If', "So", and "Because"', in M. Black (ed.), *Philosophical Analysis* (Prentice Hall, 1950), 302–18.

Sacks, M. (1997), 'Transcendental Constraints and Transcendental Features', *International Journal of Philosophical Studies* 5, 164-186.

Salmon, N. (2005), *Reference and Essence*, 2nd Edition, Prometheus Books.

Sandis, C. (2010), 'The Experimental Turn and Ordinary Language', *Essays in Philosophy* 11, 181-196.

Sandis, C. (2019), 'Who Are "We" for Wittgenstein?', in Appelqvist (ed.), *Wittgenstein and the Limits of Language*, London: Routledge, 172–196.
Sandis, C. (2021), 'No Picnic: Cavell on Rule-Descriptions', *Philosophical Investigations*, 295-317.
Sankey, H. (1994), *The Incommensurability Thesis*, Aldershot: Avebury.
Sankey, H. (1997), *Rationality, Relativism and Incommensurability*, Aldershot: Ashgate.
Sankey. H. (2011), 'Incommensurability and Theory Change', in Hales (ed.), *A Companion to Relativism*, Oxford: Wiley-Blackwell, 456-474.
Schauer, F. (1991), 'Exceptions', *The University of Chicago Law Review* 58, 871-899.
Schauer, F. (2011), 'On the Open Texture of Law', *Grazer Philosophische Studien*, 197–215.
Scheffler, I. (1967), *Science and Subjectivity*, Indianapolis: Bobbs-Merrill.
Schlick, M. (1918), *Allgemeine Erkenntnislehre*, Berlin: Springer.
Schönbaumsfeld (2010), 'A "Resolute" Later Wittgenstein?', *Metaphilosophy*, 649-668.
Schroeder, S. (2006), *Wittgenstein: The Way Out of the Fly-Bottle*, Cambridge: Polity.
Schroeder, S. (2009), 'Analytical Truths and Grammatical Propositions', in Glock and Hyman (eds.), *Wittgenstein and Analytic Philosophy: Essays for P.M.S. Hacker*, Oxford: Oxford University Press, 83-108.
Schroeder, S. (2011), 'Conjecture, Proof, and Sense in Wittgenstein's Philosophy of Mathematics', *Epistemology: Contexts, Values, Disagreement: proceedings of the 34th Ludwig Wittgenstein Symposium in Kirchberg,* 459-471.
Schroeder, S. (2013), 'Wittgenstein on Rules in Language and Mathematics', in Venturinha (ed.), *The Textual Genesis of Wittgenstein's 'Philosophical Investigations'*, London: Routledge, 155-67.
Schroeder, S. (2014), 'Mathematical Propositions as Rules of Grammar', *Grazer Philosophische Studien* 89, 23-38.
Schroeder, S. (2015), 'Mathematics and Forms of Life', *Nordic Wittgenstein Review*, Special Issue, 111-130.
Schroeder, S. (2017), 'Wittgenstein on Grammar and Grammatical Statements', in Glock and Hyman (eds.), *The Blackwell Companion to Wittgenstein*. Oxford: Wiley-Blackwell, 252-268.
Schroeder, S. (2021), *Wittgenstein on Mathematics*, Abingdon: Routledge.
Searle, J. (1969), *Speech Acts*, Cambridge: Cambridge University Press.
Sellars, W. (1954), 'Presupposing', *Philosophical Review* 63, 197–215.
Shapiro, S. (2000), 'The Status of Logic', in Boghossian and Peacocke (eds.), *New Essays on the A Priori*, Oxford: Oxford University Press, 333-366.
Shapiro, S. (2006), *Vagueness in Context*, Oxford: Oxford University Press.

Shapiro, S. (2018), 'Changing the Subject: Quine, Putnam and Waismann on Meaning-Change, Logic, and Analyticity', in Hellman and Cook (eds.), *Hilary Putnam on Logic and Mathematics*, Springer, 115-126.

Shapiro, S. and Roberts, C. (2019), 'Open Texture and Analyticity', in Makovec and Shapiro (eds.), *Friedrich Waismann: The Open Texture of Analytic Philosophy*, Palgrave Macmillan 189-210.

Sidelle, A. (1989), *Necessity, Essence and Individuation*, Ithaca: Cornell University Press.

Sidelle, A. (2002). 'On the Metaphysical Contingency of Laws of Nature', *Conceivability and Possibility*, Gendler and Hawthorne (eds.), Oxford: Oxford University Press, 309-336.

Sidelle, A. (2007), 'The Method of Verbal Dispute', *Philosophical Topics* 35, 83-113.

Sidelle, A. (2009), 'Conventionalism and the Contingency of Conventions', *Nous*, 224-241.

Sider, T. (2003), 'Reductive Theories of Modality', in Loux and Zimmerman (eds.), *The Oxford Handbook of Metaphysics*, Oxford: Oxford University Press, 180-208.

Sider, T. (2011), *Writing the Book of the World*, Oxford: Oxford University Press.

Simion, M. and Kelp, C. (2020), 'Conceptual Innovation, Function First', *Nous* 54, 985–1002.

Smith, Q. (1995), 'Marcus, Kripke, and the Origin of the New Theory of Reference', *Synthese* 104, 179-189.

Soames, S. (1989), 'Presupposition', in Gabbay and Guenthner (eds.), *Handbook of Philosophical Logic*, Dordrecht: Springer, 553-616.

Soames, S. (1999), 'The Indeterminacy of Translation and the Inscrutability of Reference', *Canadian Journal of Philosophy* 29, 321-370.

Soames, S. (2003), *Philosophical Analysis in the Twentieth Century: The Age of Meaning*, Vol. 2, Oxford: Princeton University Press.

Soames, S. (2006), 'The Philosophical Significance of The Kripkean Necessary Aposteriori', *Philosophical Issues* 16, 288-309.

Soames, S. (2007) 'The Substance and Significance of the Dispute over Two Dimensionalism' *Philosophical Books*, 34–49.

Sober, E. (2000), 'Quine's Two Dogmas', *Proceedings of the Aristotelian Society*, Supplementary 74, 237–280.

Sorensen, R. (2001) 'Vagueness has No Function in Law', *Legal Theory* 7, 387–417.

Sorensen, R. (2018), "Vagueness", *The Stanford Encyclopedia of Philosophy* (Summer 2018 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/sum2018/entries/vagueness/>.

Srinivasan, A. (2019), 'Genealogy, Epistemology and Worldmaking', *Proceedings of the Aristotelian Society* 119, 127-156.

Sterken R. K. (2020), 'Linguistic Intervention and Transformative Communicative Disruptions', in Burgess, Cappelen and Plunkett, (eds.), *Conceptual Engineering and Conceptual Ethics*, Oxford: Oxford University Press, 417-434.

Stern, D. (2011), 'Private language', in Kuusela and McGinn (eds.), *The Oxford Handbook of Wittgenstein*, Oxford: Oxford University Press, 333–50.

Strawson, P. F. (1950), 'On Referring', *Mind* 59, 320–44.

Strawson, P. F. (1952), *Introduction to Logical Theory*, London: Methuen.

Strawson, P. F. (1954), 'A Reply to Mr. Sellars', *Philosophical Review* 63, 216–23.

Strawson, P. F. (1963), 'Carnap's Views on Constructed Systems Versus Natural Languages in Analytic Philosophy', in Schilpp (ed.), *The Philosophy of Rudolf Carnap*, Lasalle: Open Court, 503-518.

Strawson, P. F. (1966), *The Bounds of Sense: An Essay on Kant's 'Critique of Pure Reason'*, London: Methuen.

Stroll, A. (1989), 'What Water Is or Back to Thales', *Midwest Studies in Philosophy*, 258-274.

Stroll, A. (1994), *Moore and Wittgenstein On Certainty*, Oxford: Oxford University Press.

Stroud, B. (1965), 'Wittgenstein and Logical Necessity', *Philosophical Review*, 504-518.

Stroud, B. (1968), 'Transcendental Arguments', *Journal of Philosophy* 65, 241-256.

Suits, B. (1978), *The Grasshopper: Games, Life and Utopia*, Broadview Press.

Sullivan, P. (2003), 'Ineffability and Nonsense', *Proceedings of the Aristotelean Society*, 195-223.

Sultanescu, O. (2022), 'Meaning, Rationality, and Guidance', *Philosophical Quarterly* 73, 227–247.

Tahko, T. (2009), 'On the Modal Content of *A Posteriori* Necessities', *Theoria*, 344-357.

Tahko, T. (2013), 'Soames's Deflationism About Modality', *Erkenntnis*, 1367-1379.

Tahko, T. (2015), 'Natural Kind Essentialism Revisited', *Mind*, 795-822.

Thomasson, A. (2009), 'Answerable and Unanswerable Questions', in *Metametaphysics*, Chalmers, Manley, and Wasserman (eds.), Oxford University Press, 444-471.

Thomasson, A. (2020a), *Norms and Necessity*, Oxford: Oxford University Press.

Thomasson, A. (2020b), 'A Pragmatic Method for Normative Conceptual Work', in Burgess, Cappelen and Plunkett, (eds.), *Conceptual Engineering and Conceptual Ethics*, Oxford: Oxford University Press, 435-458.

Thomasson, A. (2021), 'How Can We Come to Know Metaphysical Modal Truths?', *Synthese*, Supplementary 8, 2077-2106.

Topey, B. (2019), 'Linguistic Convention and Worldly Fact: Prospects for a Naturalist Theory of the A Priori', *Philosophical Studies* 176, 1725-1752.

Travis, C. (2006), *Thought's Footing*, Oxford, New York: Oxford University Press.

Tripodi, P. (2020), *Analytic Philosophy and the Later Wittgensteinian Tradition*, Palgrave MacMillan.
Vendler Z. (1967), *Lingusitics in Philosophy*, Ithaca: Cornell University Press.
von Wright, G. H. (1983), 'Norm, Truth and Logic', *Practical Reason: Philosophical Papers*, Vol. 1, Oxford: Blackwell, 130–209.
von Wright, G. H. (1993), 'A Pilgrim's Progress', *The Tree of Knowledge*, Brill Academic Publishers, 103–113.
Waismann, F. (1945), 'Verifiability', *Proceedings of the Aristotelian Society*, Supplementary Volumes, Vol. 19, Analysis and Metaphysics, 119-164.
Waismann, F. (1949), 'Analytic-Synthetic I', *Analysis*, 25-40.
Waismann, F. (1950), 'Analytic-Synthetic II', *Analysis*, 25-38.
Waismann, F. (1951a), 'Analytic-Synthetic III', *Analysis*, 49-61.
Waismann, F. (1951b), 'Analytic-Synthetic IV', *Analysis*, 115-124.
Waismann, F. (1952), 'Analytic-Synthetic V', *Analysis*, 1-14.
Waismann, F. (1953), 'Analytic-Synthetic VI', *Analysis*, 73-89.
Waismann, F. (1965), *The Principles of Linguistic Philosophy*, Harré (ed.), London: Macmillan, New York: St. Martin's Press.
Waismann, F. (1968), *How I See Philosophy*, New York: Macmillan.
Waldron, J. (1994), 'Vagueness in Law and Language: Some Philosophical Issues', *California Law Review* 82, Symposium: Void for Vagueness, 509-540.
Walker, R. C. S. (1989), *The Coherence Theory of Truth*, London: Routledge.
Warren, J. (2015a), 'The Possibility of Truth by Convention', *The Philosophical Quarterly* 65, 84–93.
Warren, J. (2015b), 'Talking with Tonkers', *Philosophers' Imprint* 15, 1–24.
Warren, J. (2017), 'Revisiting Quine on Truth by Convention', *Journal of Philosophical Logic* 46, 119–139.
Warren, J. (2018), 'Change of Logic, Change of Meaning', *Philosophy and Phenomenological Research* 96, 421–442.
Warren, J. (2020), *Shadows of Syntax: Revitalizing Logical and Mathematical Conventionalism*, New York: Oxford University Press.
Warren, J. (2022), *The A Priori Without Magic*, Cambridge: Cambridge University Press.
Weinberg, J., Nichols, S., and Stich, S. (2001), 'Normativity and Epistemic Intuitions', *Philosophical Topics* 29, 429–460.
Westphal, K. (2015), 'Conventionalism and the Impoverishment of the Space of Reasons: Carnap, Quine and Sellars', *Journal for the History of Analytical Philosophy* 3, 1-66.

Williams, B. (1981), 'Wittgenstein and Idealism', *Moral luck*, Cambridge: Cambridge University Press, 144-163.

Williams, B. (2002), *Truth and Truthfulness*, Princeton: Princeton University Press.

Williams, M. (2002), 'Review of Hacker, Wittgenstein: Connections and Controversies', *Notre Dame Philosophical Reviews*.

Williamson, T. (1994). *Vagueness*, London: Routledge

Williamson, T. (2007), *The Philosophy of Philosophy*, Oxford: Wiley-Blackwell.

Wright, C. (1986), 'Inventing Logical Necessity', in Butterfield (ed.), *Language, Mind and Logic*, Cambridge: Cambridge University Press, 187–209.

Wright, C. (2008), 'Rule-Following Without Reasons: Wittgenstein's Quietism and the Constitutive Question', in Preston (ed.), *Wittgenstein and Reason*, Oxford: Blackwell, 123-144.

Yablo, S. (1992), 'Review of Necessity, Essence, and Individuation: A Defense of Conventionalism', *Philosophical Review* 101, 878–881.

Yap, A. (2010), 'Feminism and Carnap's Principle of Tolerance', *Hypatia* 25, 437-454.