

Mark J. Cain

Two Accounts of Concept Possession

Abstract

In this paper I examine the conflict between two radically different accounts of concept possession, one due to Ludwig Wittgenstein and the other due to Jerry Fodor. That conflict centres around the viability of atomism and mentalism. Wittgenstein's rejection of atomism opens him to a version of Fodor's familiar objection to non-atomistic positions. I argue that there is little prospect of blunting the force of this objection. Moreover, on closer inspection, atomism turns out not to be as implausible as is often thought.

With respect to mentalism, Wittgenstein's frequent criticisms of mentalist theorising suggests a parallel objection to Fodor's position. The power of this objection ultimately depends upon whether concepts and rules have normative properties that preclude their being grounded in causal and mechanical phenomena. On this point I argue that there are grounds for the Fodorian to be optimistic. In the light of all this I conclude that Fodor's account of concept possession is to be preferred to Wittgenstein's.

Introduction

What is it to possess a concept? In this paper my aim is to examine the conflict between two diametrically opposed answers to this question. One of these answers is associated with the work of the later Wittgenstein whilst the other has been vigorously championed by Jerry Fodor. I will argue that the Wittgensteinian case against Fodor centres on an opposition to atomism and to mentalism and that that case is unconvincing. Moreover, Fodor's work raises a number of powerful objections to Wittgenstein's account of concept possession.

Wittgenstein on Concept Possession

Wittgenstein does not explicitly present an account of concept possession. Rather, the account that I am going to attribute to him is implicit in much of his latter work, particularly in his reflections on the nature of philosophy in the *Philosophical Investigations*. For Wittgenstein philosophy is a grammatical investigation. Conducting a grammatical investigation involves describing the use that we make of words and sentences in the linguistic practices in which we use them. Philosophy aims at description rather than explanation and it describes something that all competent language users are familiar with. Thus, philosophy does not reveal anything “that is hidden” and “leaves everything as it is” (1953: § 124). All this is perhaps best expressed by Wittgenstein at *Philosophical Investigations* § 109 where he writes:

It was true to say that our considerations could not be scientific ones. ... And we may not advance any kind of theory. There must not be anything hypothetical in our considerations. We must do away with all *explanation*, and description alone must take its place. And this description gets its light, that is to say its purpose, from the philosophical problems. These are, of course not empirical problems; they are solved, rather, by looking into the workings of our language, and that in such way as to make us recognize those workings: *in despite* of an urge to misunderstand them. The problems are solved, not by giving new information, but by arranging what we have always known. Philosophy is a battle against the bewitchment of our intelligence by means of language.

What, one might ask, is the point of such a grammatical investigation? Wittgenstein’s answer is that it can reveal the error and confusion that we perennially fall into when give into the urge philosophise and can help us to avoid making such mistakes in the future. It is difficult to see how a grammatical investigation can produce such results were there not a significant connection between linguistic usage and the nature of the concepts that we employ, a connection such that to possess a given concept is to have mastery of the use of relevant words and sentences. Wittgenstein’s commitment to such an account of concept possession is further revealed in his reflections on meaning and understanding where he attacks the idea that meaning something by a word or sentence or understanding a word or sentence in a particular way is a mental state that lies behind linguistic use. Rather, it is the use that we make of words and sentences that breathes life into otherwise dead signs. Thus, we might characterise Wittgenstein’s account of concept possession in the following terms:

To possess a concept C is to have a mastery of the use of relevant words and sentences. For example, for a member of an English speaking community, to possess the concept DOG is to have mastery of the use of the word 'dog' and sentences containing that word.

As this account makes an appeal to linguistic usage, some light needs to be shed upon that notion. When Wittgenstein talks about the use of a word or sentence he is referring to a range of phenomena such as the following. First, the act that we use the word or sentence to accomplish, for example, state a fact, name an object, ask a question, express a feeling, make an order, and so on. Second, the circumstances in which we utter the word or sentence. For example, suppose that I say 'that's an aardvark' in response to the question 'what is that?' when confronted by an aardvark. Then, describing my use of 'that's an aardvark' involves describing the surrounding circumstances of my being confronted by an aardvark and asked the question 'what is that?' Third, if a word or sentence is used to make a statement then the factors that we regard as evidence for its truth and the manner in which we go about determining its truth belong to its use as do the conclusions that we draw from it if we regard it as true. Fourth, the factors that we appeal to in order to justify or criticise the utterance of a word or sentence belong to its use.

Thus, Wittgenstein's notion of use is an inherently intentional one; he is not involved in any project of reducing concept possession to something non-intentional, for example, physically characterised behaviour.¹

An important implication of Wittgenstein's notion of use is a commitment to the rejection of an atomistic view of concepts, the view that the content of a concept is not determined by its relations to other concepts. This comes out in Wittgenstein's remark that 'to understand a sentence means to understand a language.' (1953: § 199). The idea here is that the uses of distinct words and sentences are intimately bound up with one another. For example, it is part of the use of the sentence 'Alf is an aardvark' that the truth of 'Alf eats termites' is taken as grounds for asserting it and that its truth is grounds for asserting that 'Alf is an animal.' Here there is a strong analogy with chess. The use of the knight is intimately bound up with that of all the other pieces so that it is not possible to have mastery of the use of the knight without having mastery of the use of all the other pieces. The implication with respect to concept possession is that, for example, one cannot possess the concept AARDVARK without possessing a whole battery of other concepts such as EATS TERMITES and ANIMAL.

¹ This distinguishes Wittgenstein from most contemporary advocates of Conceptual Role Semantics whose theories are sometimes described as use theories. See, for example, Block (1986) Harman (1982) and Loar (1981).

Another important point about linguistic use as conceived by Wittgenstein is brought out in his comparison of language and games as captured in the notion of a language game. Playing games involves doing things, for example, moving pieces in chess, kicking a ball in football. Such activities are observable; they involve behaving or acting and one's mastery of the game is manifested in such behaviour. Having mastered the game involves being able to behave in appropriate ways in circumstances thrown up in the course of a game. Language is likewise a spatial and temporal phenomenon (1953: § 108) in the respect that it is bound up with nonlinguistic behaviour and activities. As he says in § 23 'the term "language-game" is meant to bring into prominence the fact that speaking a language is part of an activity, or a form of life.' This comes out in the case of the simple language game of the builders as described in § 1. The assistant's mastery of the language game involves his responding to a call of 'slab' by passing a slab to the builder, 'beam' by passing a beam, and so on. Consequently, using a word or sentence often involves interacting with the extra linguistic world and engaging in non-linguistic behaviour.²

Not everything that a person might do with an expression belongs to its use. For example, suppose I say of an aardvark that I confront on a dark night 'that's an anteater'. Or suppose I make an assertion with a sentence on the grounds of my acceptance of some other sentences where the truth of the latter sentences would not be ordinarily regarded as grounds for committing oneself to the truth of the former. In such cases I have misused the sentence in question. Such 'uses' do not belong to the use any more than weighing down paper belongs to the function of a hammer just because I sometimes use my hammer as a paperweight. 'Use' means something like 'proper use'; what one is supposed to do with the sentence rather than what one actually does. This connects use to rules and therefore gives it a normative dimension. There are rules governing the use of the expressions of a language and certain things we do with an expression violate these rules and are thus misapplications or misuses of the expression. Therefore, to possess a concept is to have mastered a rule governed practice.

Fodor on Concept Possession

Fodor is explicit on what it is to possess a concept and his account can be expressed in the following terms. To have a concept is to have a mental representation in one's Language of Thought (LOT for short) that has the

² The world involving character of use as conceived by Wittgenstein serves to distinguish his position from that of Brandom (1994, 2000).

appropriate content. For example, to have the concept AARDVARK is to have a symbol in one's LOT that means *aardvark*. Content reduces to information: it is a matter of mind-world causal or nomic relations.³ So, to a first approximation, a mental representation means *aardvark* if and only if its tokenings are reliably caused by aardvarks. Hence, conceptual content is atomistic so that, at least in principle, one could have the concept AARDVARK without having any other concepts. As Fodor puts it 'satisfying the metaphysically necessary conditions for having one concept never requires satisfying the metaphysically necessary conditions for having any other concept' (1998: 13-14). But not everything that means *aardvark* is a concept and a system (say, a machine or an animal) could have an internal state with a given content without having the corresponding concept. As, for something to be a representation belonging to LOT, it must have the right role within an internal economy of considerable complexity; it must be available to be manipulated by appropriate inhead computational mechanisms.⁴

There is much in Fodor's position that Wittgenstein would have found deeply problematic and to fully appreciate this it will be necessary to examine two elements of the broader context in which Fodor's account is placed. First, Fodor's account is overtly mentalist as he conceives of the mind as being an inner system within which mental phenomena reside. In general, mental phenomena are token identical to, or constituted by, neural phenomena and lie behind, typically causing, publicly observable behaviour. Concepts, being symbols of LOT, reside within the mind and, accordingly, are neurally embodied in being token identical to, or constituted by, brain states.⁵ For Fodor the connection between concept possession and linguistic mastery is indirect. In principle, an individual could have a rich conceptual scheme the elements of which she frequently employed in thought without having a language. For language users, the purpose of language is to enable them to communicate their thoughts, thoughts that they have independently of their language. Insofar as an individual's utterances have semantic properties, those properties are inherited from the underlying intentional states that cause the utterances, states that in turn inherit their content from the symbols of LOT that they feature. Therefore, at best, an individual's use of language provides indirect evidence of the concepts that she possesses and the intentional states that she tokens.

³ This needs to be qualified as Fodor holds an informational theory of content only with respect to simple, non-logical symbols of LOT and not with respect to logical symbols and names. See Fodor (1990a).

⁴ See Cain (2002) for detailed discussion.

⁵ For a clear account of the distinction between token identity and constitution see Rudder Baker (1995).

Wittgenstein's account of concept possession is overtly non-mentalist in that it does not represent the concepts that an individual possesses as being a matter of how things are in her mind conceived as an inner system. This reflects one of the dominant themes in Wittgenstein's later writings as, throughout his reflections on the Augustinian picture of language, rule-following and the possibility of a private language, he sought to reveal the explanatory bankruptcy of mentalism.

Second, Fodor's account of concept possession is part of a general project the aim of which is to explain familiar facts about us. According to Fodor, intentional states such as beliefs and desires are computational relations to LOT symbols and intentional processes such as reasoning are computational processes involving the manipulation of LOT symbols. One of Fodor's chief arguments for this theory is that it explains a series of features of our mental lives that cry out for explanation, namely that thought is productive, systematic, intensional and rationally coherent.⁶ Such an explanation is akin to a scientific explanation in that it explains a prominent observable feature of the world by postulating unobserved phenomena. In fact, there is no fundamental distinction between philosophy and science and philosophy of mind done properly is a branch of cognitive science. Moreover, in *The Language of Thought* Fodor characterised his postulation of LOT as an exercise in speculative psychology that was supported by the best psychological theories of cognition of the day.⁷ Thus, Fodor is at loggerheads with Wittgenstein's antiexplanatory conception of philosophy that represents philosophy as a descriptive enterprise that is not concerned with anything that is hidden and is to be sharply distinguished from science.

Which of these two divergent accounts of concept possession is the most plausible? In the remainder of this paper I will argue in favour of Fodor's account. I will do this by examining the conflict that arises between Wittgenstein and Fodor over the issues of atomism and mentalism.

Concept Atomism

Atomism is the doctrine that the content of a concept is not a matter of its relations to other concepts. Atomism is to be contrasted with holism, the view that the content of a concept is a matter of the relations that it bears to all other concepts

⁶ See Fodor (1957: Appendix) for the clearest statement of this argument.

⁷ More recently, cognitive theories opposed to the LOT model have been developed, particularly within the connectionist framework. Fodor has responded to such theories by highlighting their explanatory limitations. In particular, he has argued that connectionism cannot account for the systematicity of thought. See Fodor and Pylyshyn (1998) and Fodor and McLaughlin (1990).

in the conceptual scheme to which it belongs.⁸ In endorsing an informational theory of content Fodor is explicitly committed to concept atomism. We have seen that, in contrast, Wittgenstein implicitly rejects atomism as the use of any given word or sentence is bound up with that of many others. Many philosophers find atomism a thoroughly unappealing doctrine so Wittgenstein might be seen to have the advantage here.⁹ However, that would be to overlook a problem that, as Fodor has repeatedly emphasised, plagues non-atomistic views of content. The problem is this: if one rejects atomism then one makes it a complete mystery how, barring an accident on a cosmic scale, any two individuals or time slices of the same individual can share a concept. This problem generates an objection to the effect that non-atomism is inconsistent with a perfectly familiar fact, namely, the fact that distinct individuals, and time slices of the same individual, share many of their concepts.¹⁰ In greater detail and modified so as to specifically to apply to Wittgenstein the argument goes as follows.

It would appear that two individuals can share a concept and express that concept by means of one and the same word yet have conflicting beliefs involving that concept. Consider an example. Both you and I grasp the concept AARDVARK and express that concept by means of the English word 'aardvark'. However, our respective beliefs about aardvarks differ considerably. You believe that aardvarks are water loving and practically aquatic whilst I believe that they are fearful of water and are loath to go swimming. These differences between our respective beliefs about aardvarks are such as to influence the use that we make of the word 'aardvark' and many sentences containing that word so that we use those words and sentences differently. To see this consider our respective use of the sentence 'that is an aardvark' as used to categorise a seen object. You are in circumstances where you see an animal with an external appearance like that of an aardvark contentedly immersed in a pool of water and are asked what the creature is. Believing that aardvarks are water loving, you respond by saying 'that is an aardvark.' In the same circumstances I, believing that aardvarks are fearful of water, do not respond by saying 'that is an aardvark.' I reason that as aardvarks are fearful of water the animal contentedly immersed in a pool of water can't be an aardvark even though it looks quite like one. This is a significant difference in use and there are other such differences generated by our difference

⁸ Lying in between the extremes of atomism and holism is molecularism, the view that the content of a concept is a matter of the relations that it bears to some, but not all, of the concepts that belong to the conceptual scheme that it belongs to. I shall use the term non-atomism to refer to a position that covers both molecularism and holism.

⁹ The list of Philosophers who reject atomism makes impressive reading and includes Sellars (1956), Quine (1960), Dennett (1987), Davidson (1984), Lewis (1983), and Churchland (1991).

¹⁰ See Fodor (1987: Ch 3) and Fodor and LePore (1992).

of opinions on the question of the attitude of aardvarks to water. You are prone to regard ‘that animal is contentedly immersed in a pool of water’ as grounds for saying ‘that is an aardvark’ and are prone to infer from the latter sentence the sentence ‘that animal is fond of water.’ But with respect to me matters are quite the opposite. In short there are differences in the respective uses that we make of ‘that is an aardvark.’ For the Wittgensteinian this implies a difference in concepts where, *ex hypothesi*, there is none. If the Wittgensteinian bites the bullet and concludes that you and I do not express the same concept by means of the word ‘aardvark’ then she seems to be landed with a horrific consequence. Take any two individuals or time slices of the same individual. They are bound to differ in their beliefs in such a way as to give rise to differences in the use that they make of any given word or sentence that belongs to both of their languages. Consequently, they do not share any concepts. Quite apart from the inherent implausibility of this consequence, it doesn’t sit very happily with Wittgenstein’s opposition to the possibility of a private language and his portrayal of humans as belonging to communities sharing a common language.

One obvious line of response to this objection would be to argue that not everything that belongs to the putative use of a word or sentence belongs to its actual use. In the case of the present example it might be argued that inferring ‘that is an animal’ from ‘that is an aardvark’ is part of the use of ‘aardvark’ whereas inferring ‘it is fond of water’ from ‘that is an aardvark is not.’ The problem with this line of thought is that a substantial account is needed of what distinguishes those elements of putative use that belong to actual use from those that do not. One traditional way of doing this involves an appeal to the analytic- synthetic distinction. So, the idea might go, as ‘all aardvarks are animals’ is analytically true to master the use of ‘aardvark’ one has to be disposed to infer ‘x is an animal’ from ‘x is an aardvark.’ However, as ‘all aardvarks are fond of water’ is not analytically true one can master the use of ‘aardvark’ without being remotely disposed to infer ‘x is fond of water’ from ‘x is an aardvark.’

There are several problems with this idea. First, since Quine (1951) it has become clear that it is far from obvious that there is an analytic-synthetic distinction. Moreover, it may well be that Wittgenstein himself is committed to rejecting the distinction in arguing that our words rarely can be defined and that our concepts typically do not have necessary and sufficient conditions for their application.¹¹

Second, even if there is an analytic-synthetic distinction, it is far from clear that there are the required analytic truths corresponding to all of our concepts. Take the classic example of ‘bachelor.’ Not only is it analytically true that

¹¹ See Wittgenstein (1953: §§ 66-77).

'bachelors are unmarried men' but that sentence specifies what it is to be a bachelor. But there doesn't appear to be a corresponding analytic truth with respect to 'aardvark.' It might be analytically true that 'aardvarks are animals' but that sentence doesn't exhaustively tell us what it is to be an aardvark. Any sentence that did exhaustively tell us what it is to be an aardvark would be a poor candidate for an analytically true sentence and would only be known by a small number of biologists. So one could say that to master 'aardvark' one has to be disposed to infer 'x is an animal' from 'x is an aardvark.' But for the use theory to be plausible there has to be more to the use of 'aardvark' than this as it is also analytically true that 'anteaters are animals.' The worry is that unless there are a whole battery of analytic truths featuring 'aardvark' that serve to distinguish it from all other nonsynonymous words of English then not enough belongs to the use of 'aardvark' for that word to express the concept AARDVARK. The same would hold, *mutatis mutandis*, of most other words of English.

Third, there is the worry that appealing to the analytic-synthetic distinction implies that in the normal course of things we do not know what our fellows mean by their words and sentences or what concepts they possess. Suppose that I am trying to determine how you use 'aardvark' and whether you use that word as I do. Gathering evidence concerning the inferences that you draw involving sentences containing 'aardvark' will not help me all that much as I have to determine which of these inferences correspond to analytic truths featuring your word 'aardvark' and which do not. But how am I to do this? The worry is that I have to independently know what you mean by 'aardvark' before I can discover which of the inferences that you make are relevant to the meaning of that word or the identity of the concept that you express by it. It is not as if I could ask you which sentences you regard as analytically true as even those individuals who understand such questions are often unsure on such matters. Is 'aardvarks are animals' analytically true? I'm not sure but I think it has a better claim to being analytic than 'aardvarks eat termites.' Suppose that I was sure on such matters, could I be mistaken? I think so as there are examples that suggest that individuals are sometimes not authoritative one whether sentences they hold as true are analytic or synthetic. For example, when I ask my students whether ' $2+2=4$ ' is analytic or synthetic they tend to split down the middle. Assuming that my students share their mathematical concepts, half of them are mistaken on this matter. If, as I have suggested, it is going to be difficult to determine which elements of an individual's use of a given word or sentence correspond to analytic truths and which are the product of beliefs that could be abandoned without any change in meaning or conceptual content, then it is going to be difficult to determine what any individual means by their words and sentences and which concepts they possess. That is not to say that there is no fact of the matter but

only that we have problems acquiring knowledge of what the facts of the matter are in this area. This would be a problem for Wittgenstein as he is keen to emphasise that meanings and concepts are typically shared. It would be no good to make an appeal to the community at this point by saying that it is the use of the community that the individual belongs to that determines her individual use. For one thing, an individual's use of a word or sentence could diverge from that of the community to which she belongs.¹² For another, there are going to be problems determining the community's view on such matters and working out when a putative change in usage reflects a change in meaning or conceptual content and when it is the product of a change in contingently held beliefs.

In sum then, there are substantial reasons for doubting than an appeal to analyticity could help the Wittgensteinian to defeat the Fodorian objection that Wittgenstein's account of concept possession implies that no two people or time slices of the same person could share a concept. An alternative response to that objection would be to argue that there are similarities of use that engender sameness of content so that two individuals can share a concept if there are appropriate similarities between their respective use of a given word and sentences that contain it. The problem with attempting to make this move is that one is thereby given a substantial additional burden of providing a disciplined and motivated account of what similarity of use comes to. What is it about your and my respective use of 'aardvark' that makes that use similar in a way in which my use of 'aardvark' is not similar to my use of 'anteater'? Perhaps there is a viable answer to this question but it needs to be established that there is such an answer and what it is. For the Wittgensteinian, the immediate worry is that such a burden could not be met without adopting the kind of theoretical approach that would be judged illegitimate by Wittgenstein. It is not enough to say that our respective use of 'aardvark' is similar because members of our linguistic community would apply the word 'similar' to those uses, this being part of their use of 'similar.' This is because the problem that the notion of similarity is being invoked to solve arises in connection with the use of 'similar'.

In the light of these problems it might be worthwhile reassessing the credentials of atomism. Many philosophers find atomism implausible as it appears to grant the possibility of a creature's having only one concept. How, they might contend, could a creature be capable of thinking that the objects that it encountered were aardvarks without being able to think of them in any other way? A related worry is that the atomist makes it too easy to possess a concept by implying that any system (including simple animals and machines) that discriminatively

¹² To say this is not to commit oneself to the possibility of a private language of the kind that Wittgenstein famously discusses.

responds to a given property thereby has the concept of that property. In response it should be pointed out that Fodor's account of concept possession has no such implications as his central claim is that concepts are symbols of LOT. For the tokening of an internal state within a system to count as the application of the concept AARDVARK, it is not enough that that state has the content *aardvark*. In addition, it must be a symbol of a LOT. And this requires the system to be capable of tokening internal states that have other contents and of combining these states to form more complex internally structured content bearing states.

It might be objected that an implicit rejection of atomism underlies the methods that we employ in determining whether an individual has any given concept. Suppose we want to know whether a particular individual has the concept AARDVARK. Typically, we would ask them a question such as 'what is an aardvark?' and if their answer suggested that they did not believe aardvarks to be animals, to be quadrupeds, to eat termites and so on, then we would not conclude that they possessed the concept AARDVARK.

In this context, it is very important appreciate that Fodor can consistently accept that many of the mind-world casual connections that he thinks determine the content of our concepts are mediated by theories or bodies of beliefs.¹³ Similarly, he can concede that, as a matter of fact, to acquire certain concepts we must form a body of beliefs as to the nature of the property expressed by the target concept. For example, suppose that in order to acquire the concept AARDVARK a human individual has to form a body of beliefs or a theory as to the salient characteristics of aardvarks, otherwise she would not be remotely disposed to respond to aardvarks by tokening the LOT analogue of 'aardvark.'. It is just that the precise nature of this body of beliefs is irrelevant to the content of that symbol; in principle two individuals could have radically different theories about aardvarks but still share the concept AARDVARK, as their respective theories both served to mediate a causal connection between aardvarks and tokenings of AARDVARK. Once this is appreciated, it becomes clear that the atomist is in a good position to deal with our prima facie non-atomistic intuitions and practices. First, it is consistent with atomism that one finds it implausible that a person could acquire any given concept without having a prior grasp of a whole battery of other concepts. For, if they don't have a prior grasp of a whole battery of other concepts, how could they form a body of beliefs that could serve to mediate a causal connection between the property expressed by the target concept and applications of that concept? Second, it is consistent with the truth of atomism that a good way of determining whether a person grasps a particular concept involves determining whether she has true beliefs about the property

¹³ See Fodor (1994: Ch 2).

expressed by the concept. For example, suppose that an individual asserts that she believes that 'aardvarks are quadrupedal termite eating animals'. Then she has provided evidence that her beliefs are such as to mediate a causal connection between aardvarks and her tokenings of AARDVARK. If, on the other hand, she appears not even to believe that aardvarks are animals then, assuming that she grasps the concept ANIMAL, we have evidence that her beliefs are not such as to mediate an aardvark-AARDVARK connection. For, if a person who grasps the concept ANIMAL doesn't believe that aardvarks are animals, then she is hardly likely to respond to an encounter with an aardvark by tokening AARDVARK.

Another point on the *prima facie* intuitive implausibility of atomism is that there are cases where the non-atomist is committed to the claim that two individuals fail to share a concept and in making this commitment clashes with common-sense. Hence, it is far from obvious that nonatomism has any clear-cut advantages over atomism in terms of its intuitive plausibility. Fodor (1987) presents an example of such a case. The ancient Greeks had a theory about stars very much at odds with that widely held today. For they 'believed that stars are little holes in the sky which the heavenly fires show through' (88). Consequently, the relation that the ancient Greeks concept STAR bore to their other concepts is very much at odds with that that our concept STAR bears to the other concepts that we hold. For example, I am not remotely disposed to infer the belief that *x* is a hole in the sky from the belief that *x* is a star. Given this divergence of inferential role, the non-atomist is committed to the idea that the ancient Greeks did not share our concept STAR. But this is clearly at odds with common-sense intuition. For, common-sense has it that the ancient Greeks shared our concept STAR but differed from us in that they had a whole collection of beliefs about stars that we do not hold and which we take to be false. If this were not the case, then we would have no hope of understanding or characterising the beliefs that the ancient Greeks expressed by using the word that is standardly translated as 'star'. Indeed, that translation would be illegitimate.

That completes my discussion of the conflict between Wittgenstein and Fodor regarding atomism. In sum, Wittgenstein faces a major problem and it is difficult to see how any appeal to analyticity or similarity could help him overcome this problem. Moreover, on closer examination, atomism is nowhere near as implausible as it is often taken to be.

Mentalism

Fodor's account of concept possession is overtly mentalist and one of Wittgenstein's major goals was to undermine mentalism. Wittgenstein held that

there is a very strong temptation to postulate inner mental phenomena when attempting to explain familiar cognitive and behavioural phenomena. He held that we should resist this temptation as mentalist explanations are explanatorily bankrupt in that they are either circular or lead to infinite regress. This is because the inner phenomena that the mentalist postulates are akin to those that they are invoked to explain. Consequently, the postulated inner phenomena are as much in need of explanation as the target phenomena that they are supposed to explain so that the mentalist has made no progress at all.¹⁴ In this section my concern is to answer the question of whether this line of thought serves to undermine Fodor's position.

Fodor's account of concept possession is part of a broader theory according to which intentional states are computational relations to sentences of LOT and intentional processes are computational processes involving the manipulation of sentences of LOT. He argues that this theory explains a cluster of properties of thought, something that none of its competitors can do. In particular, it explains the productivity, systematicity and intensionality of thought along with the fact that thought processes are generally rationally coherent. If these explanations are to work then the neurally embodied symbols of LOT must literally have content. It is here that a version of Wittgenstein's anti-mentalist argument comes into play. To portray the brain as manipulating content bearing symbols is, in effect, to postulate an inner intelligent agent, a homunculus. But the capacities of the homunculus are as much in need of explanation as those of the human host that they are invoked to explain. If the capacities of the homunculus are in turn explained by appeal to further homunculi then an infinite regress looms. If, on the other hand, those capacities can be regarded as brute so as not to need explanation, then so can the target human capacities so obviating any need for Fodor's postulation of LOT.

A first point in defence of Fodor is this. There are many perfectly good explanations that explain a capacity by appeal to the possession or exercise of analogous states and capacities. For example, I have an ability to play chess. This involves an ability to make legitimate moves when playing chess and recognise the legality or otherwise of any move made by my opponent even when I have never encountered that move before. A good explanation of this ability will appeal to my knowledge of the rules of the game and my ability to work out the implications of those rules in situations thrown up in the course of a game. So, one might ask, what is wrong with explaining the productivity, systematicity, intensionality and rational coherence of thought by appeal to the

¹⁴ For a prominent example of this of kind reasoning see Wittgenstein's treatment of the appeal to interpretations in the rule following considerations (1953: §201).

combination and manipulation of content bearing symbols of a neurally embodied language? The Wittgensteinian needs an answer to this question that doesn't imply the illegitimacy of the familiar explanation of my chess playing capacities.

Second, the anti-explanatory strand in Wittgenstein's thought needs to be provided with a justification, something that neither he nor his followers do. Wittgenstein was often far too quick in judging phenomena surrounding language acquisition and mastery to be brute and in pronouncing that explanation comes to an end somewhere. Where he does gesture towards explanation, his explanations are typically simplistic and empirically unmotivated. Consider his views on language acquisition. Wittgenstein objects to the overly intellectual process described by Augustine that, in effect, portrays first language acquisition as being akin to learning a second language.¹⁵ In its place Wittgenstein offers the following account. Acquiring one's first language is a gradual process of gaining a mastery of practices that involve the use of language (so called language games). We are inducted into language games by our 'teachers' by means of encouragement, example, correction, explanation, questioning, and the like. Such training works because there are natural human responses to it. That we respond in the way that we do is a brute fact, part of 'natural history' or distinctive human 'form of life' that neither has nor requires explanation. Learning a language like English is a matter of gradually being inducted into a large battery of diverse language games.

Does this explanation of language acquisition have any plausibility? It is certainly not based upon any rigorous empirical considerations and there are reasons for suspecting that it is empirically implausible. In the case of learning how to ride a bike or play chess there are considerable grounds for believing that we are actively encouraged to develop a particular skill, that our teachers adopt certain techniques and that there are standard ways of responding to the use of these techniques. But with respect to language learning there is not the same kind of evidence. For example, as we have learnt from Chomsky (1965), (1986), it just isn't the case that we go through any explicit process of training when it comes to acquiring language. Children acquire language regardless of the extent of explicit training that they receive and regardless of such factors as intelligence, social background, and so on.

Wittgenstein provides examples of language games, both real and imaginary. Yet despite exhorting the philosopher to describe, he provides little by way of explicit description of anything more than a small number of very simple imaginary language games such as that of the builders in *Philosophical Investigations* § 2.

¹⁵ He writes: 'Augustine describes the learning of human language as if the child came into a strange country and did not understand the language of the country; that is, as if it already had a language, only not this one.' (1953: § 32).

There is a grave danger that by focussing on such simple imaginary language games he grossly distorts what is involved in mastering a language and leads us to erroneously accept an explanation of how the builders acquired a mastery of their language game as an explanation of how we acquire language.

A third point in defence of Fodor against the Wittgensteinian charge of explanatory bankruptcy is this. Fodor conceives the processes of constructing and manipulating complex symbols of LOT as being entirely mechanical. They are the kinds of processes that can be executed by a machine and do not require intelligence. Moreover, the contents of the symbols so constructed and manipulated have a naturalistic, non-intentional explanation. In particular, the content of the simple non-logical symbols of LOT is determined by their mind- world causal relations. Consequently, Fodor does not postulate any homunculi and is innocent of the Wittgensteinian charge.

How should the Wittgensteinian respond to this defence of Fodor? Answering this question will introduce us to an argument that threatens to do more damage to Fodor than the simple homunculus argument described above. This argument appeals to normativity and will be the topic of the next section.

Normativity

Central to Wittgenstein's conception of language and concept possession and application is the notion of normativity. For Wittgenstein, the use of linguistic symbols and concepts is rulegoverned so that some of the things one might do with a word or concept are correct, others incorrect, some one ought to do whilst others one ought not do. A consequence of this is that the inner processes postulated by Fodor would have to have a normative dimension to count as processes of constructing and manipulating content bearing symbols. But could they have such a normative dimension? Fodor characterises the processes of constructing and manipulating symbols of LOT as being causal and mechanical and the contents of such symbols as being determined by their causal relations to phenomena in the external world. The Wittgensteinian would object that causal phenomena are not capable of grounding normative phenomena for the following reason. Causal phenomena and causal regularities are a matter of what actually happens in the world or of what would happen in various possible circumstances. Normative phenomena, in contrast, are a matter of what ought or should happen. What does or would happen is independent of what ought or should happen and one cannot infer anything about the latter from a description of the former.¹⁶ In

¹⁶ See Kripke (1980) for a vivid expression of this line of thought.

short, Fodor illegitimately attempts to reduce something that is inherently normative to something inherently non-normative. Whatever is going on in our brains when we apply concepts and engage in reasoning, it cannot be described as involving the construction and manipulation of content bearing symbols.

What are we to make of this argument? It is to be taken very seriously but as it stands it is far from decisive for two reasons. First, it is far from obvious that normativity cannot be explained in naturalistic terms. Indeed, Fodor has expended much effort on developing a naturalistic theory of content that accounts for misrepresentation and so explains the difference between correct and incorrect applications of a concept. Whatever specific problems that there are with Fodor's theory, the number of prominent and influential naturalists implies that the naturalist approach is not obviously a non-starter.¹⁷ What is needed is a powerful argument for the conclusion that there is an unbridgeable is-ought gap. The problem is that it is difficult to discern any such argument in the Wittgensteinian literature; all that one finds is an unjustified assumption.

Second, it is possible to reject the assumption that content and rules are inherently normative.¹⁸ I do not wish to deny that there are normative facts associated with many of the content bearing symbols and rules that philosophers have focussed their attention on. But such facts might well be grounded in facts concerning human social practises that need not be present in every case where a content bearing symbol is involved or where rules are operative. Consider the example of the rules of chess. In teaching a child how to play chess we typically explicitly state the rules of the game, encourage them to consciously consult the rule when deciding how to behave and aim to instil in them a respect for the rules and a desire to act in accord with them. When they violate the rules we criticise them and advise them to consider the rules and think again. And when we realise we have violated a rule we reproach ourselves, regret our actions and, even, feel shame. My idea is that normative facts surrounding chess are the products of such practices of teaching and evaluating and the desires, regrets and states of shame that they give rise to. Thus, in principle, such normative facts need not be present in all cases of rule governed behaviour. More specifically, they need not be present when behaviour comes to be governed by a particular rule without any antecedent process akin to that involved in cases of learning the rules of chess.

Now the crucial point is that Fodor does not envisage LOT or the computational rules applied in manipulating its symbols as being learnt in the

¹⁷ Prominent naturalists include Millikan (1984, 1989), Dretske (1981, 1986), Papineau (1993) and Stalnaker (1984).

¹⁸ Such a move is not unprecedented. See, for example, Horwich (1998).

way that the vocabulary of English or the rules of chess are learnt by English speakers or chess players. Rather, LOT and the rules for manipulating its symbols are innate. In short, there is an explanation for the absence of normative facts surrounding LOT and the manipulation of its symbols and it is a mistake to generalise on the basis of a consideration of symbols and rules that learnt in a social context.¹⁹

Perhaps it is going too far to deny that content and rules have any inherent normative properties. For, if a symbol has a particular content then some possible applications of the symbol are thereby going to be correct and others incorrect. And if someone's behaviour is governed by a particular rule then some of the things they might do will be in accord with the rule and others in violation of the rule. Nevertheless, this concession doesn't completely undermine the reflections of the previous two paragraphs. This is because those reflections suggest that there may well be less to the inherent normativity of content and rules than the Wittgensteinian assumes so that there is much smaller gap for the Fodorian to bridge in accounting for the inherent normative properties of the symbols of LOT and the rules governing their manipulation. Suppose that we discovered a naturalistic property that distinguished the putatively correct applications of an LOT symbol from the putatively incorrect applications. Or suppose we discovered a naturalistic property that distinguished those manipulations of an LOT symbol that were in accord with the putative rule governing its manipulation from those that violated it. Then, the less there is to the inherent normativity of content and rules the more grounds we have to be confident that such naturalistic properties are capable of generating such inherent normativity. This idea could do with much by way of defence and elaboration, something that I cannot provide in this context. But what I have presented does suggest that as it stands the Wittgensteinian appeal to normativity is hardly decisive against Fodor.

Conclusion

In this paper I have examined the conflict between two radically different accounts of concept possession. That conflict centres around the viability of atomism and mentalism, two doctrines that Fodor is committed to and Wittgenstein rejects. Wittgenstein's rejection of atomism in tying concept possession to mastery of

¹⁹ No doubt many Wittgensteinians would object to this kind of nativism but I won't consider their objections in this paper. For an argument in favour of the claim that most of our concepts are innate see Cain (forthcoming).

linguistic use opens him to a version of Fodor's familiar objection to non-atomistic positions: that it implies that, barring an accident on a cosmic scale, no two individuals or time slices of the same individual can share a concept. There is little prospect of blunting the force of this argument by appeal to analyticity or similarity. Moreover, on closer inspection, atomism turns out not to be as implausible as is often thought. With respect to mentalism, Wittgenstein's frequent criticisms of mentalist theorising suggest a parallel objection to Fodor's postulation of LOT. Fodor conceives of the processes of constructing and manipulating symbols of LOT as being mechanical and causal, and holds that the contents of these symbols are determined by their causal relations to external phenomena. Consequently, the power of the Wittgensteinian argument against Fodor will ultimately depend upon whether concepts and rules have normative properties that preclude their being grounded in causal and mechanical phenomena. On this point I have suggested that there are grounds for the Fodorian to be optimistic. In the light of all this I conclude that Fodor's account of concept possession is to be preferred to Wittgenstein's.

References

- Block, N. 1986. Advertisement for a Semantics in Psychology. *Midwest Studies in Philosophy*, 10: 257-274.
- Brandom, R. 1994. *Making it Explicit*. Cambridge, MA: Harvard University Press.
- Brandom, R. 2001. *Articulating Reasons: An Introduction to Inferentialism*. Cambridge, MA: Harvard University Press.
- Cain, M. J. 2002. *Fodor: Language, Mind and Philosophy*. Cambridge: Polity.
- Cain, M. J. Forthcoming. 'The Return of the Nativist.' *Philosophical Explorations*.
- Chomsky, N. 1965. *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.
- Chomsky, N. 1986. *Knowledge of Language: Its Nature, Origin, and Use*. New York: Praeger.
- Churchland, P. M. 1991. *A Neurocomputational Perspective*. Cambridge, MA: MIT Press.
- Davidson, D. 1984. *Inquiries Into Truth and Interpretation*. Oxford: Oxford University Press.
- Dennett, D. C. 1987. *The Intentional Stance*. Cambridge MA: MIT Press.
- Dretske, F. I. 1981. *Knowledge and the Flow of information*. Cambridge, MA: MIT Press.
- Dretske, F. I. 1986. 'Misrepresentation'. In R. Bogdan ed. *Belief*. Oxford: Oxford University Press.
- Fodor, J. A. 1975. *The Language of Thought*. New York: Thomas Y. Crowell.
- Fodor, J. A. 1987. *Psychosemantics*. Cambridge, MA: MIT Press.
- Fodor, J. A. 1990a. 'A Theory of Content, I: The Problem'. In Fodor (1990b).
- Fodor, J. A. 1990b. *A Theory of Content and Other Essays*. Cambridge, MA: MIT Press.

- Fodor, J. A. 1994. *The Elm and the Expert*. Cambridge, MA: MIT Press.
- Fodor, J. A. 1998. *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.
- Fodor, J. A. and LePore, E. 1992. *Holism: A Shopper's Guide*. Oxford: Basil Blackwell.
- Fodor, J. A. and McLaughlin, B. 1990. 'Connectionism and the Problem of Systematicity: why Smolensky's solution Doesn't Work.' *Cognition*, 35: 183-204.
- Fodor, J. A. and Pylyshyn, Z. 1988. 'Connectionism and Cognitive Architecture: A Critical Analysis'. *Cognition* 28: 3-71.
- Harman, G. 1982. 'Conceptual Role Semantics', *Notre Dame Journal of Formal Logic* 23: 242-256.
- Horwich, P. 1998. *Meaning*. Oxford: Oxford University Press.
- Kripke, S. 1982. *Wittgenstein: On Rules and Private Language*. Oxford: Basil Blackwell.
- Lewis, D. 1983. *Philosophical Papers*, vol. I. Cambridge: Cambridge University Press.
- Loar, B. 1981. *Mind and Meaning*. Cambridge: Cambridge University Press.
- Millikan, R. G. 1984. *Language, Thought and Other Biological Categories*. Cambridge, MA: MIT Press.
- Millikan, R. G. 1989. 'Biosemantics'. *Journal of Philosophy* 86: 281-297.
- Papineau, D. 1993. *Philosophical Naturalism*. Oxford: Basil Blackwell.
- Quine, W. V. O. 1951. 'Two Dogmas of Empiricism', *Philosophical Review* 60: 20-43. Reprinted in Quine (1953).
- Quine, W. V. O. 1953. *From a Logical Point of View*. Cambridge, MA: Harvard University Press.
- Quine, W. V. O. 1960. *Word and Object*. Cambridge, MA: MIT Press.
- Rudder Baker, L. 1995. *Explaining Attitudes: A Practical Approach to the Mind*. Cambridge: Cambridge University Press.
- Sellars, W. 1956. 'Empiricism and the Philosophy of Mind.' In H. Feigl and M. Scriven eds. *The Foundations of Science and the Concepts of Psychology and Psychoanalysis*, Minnesota Studies in the Philosophy of Science, volume 1. Minneapolis, MI: University of Minnesota Press.
- Stalnaker, R. 1984. *Inquiry*. Cambridge, MA: MIT Press.
- Wittgenstein, L. 1953. *Philosophical Investigations*, trans. G.E.M. Anscombe. Oxford: Basil Blackwell.