

Cain, M (2013)

Learning, concept acquisition and psychological essentialism.

Review of Philosophy and Psychology, 4 (4). pp. 577-598.

10.1007/s13164-013-0153-4

This version is available: <http://radar.brookes.ac.uk/radar/items/605da059-1bf6-0369-0eb0-127f6bde1c85/1/>

Available on RADAR: June 2014

Copyright © and Moral Rights are retained by the author(s) and/ or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This item cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder(s). The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holders.

This document is the author's final version of the journal article. Some differences between the published version and this version may remain and you are advised to consult the published version if you wish to cite from it.

Learning, Concept Acquisition and Psychological Essentialism

M.J. Cain

Abstract

In this article I will evaluate the popular view that we acquire most of our concepts by means of learning. I will do this through an examination of Jerry Fodor's dissenting views and those of some of his most persistent and significant critics. Although I will be critical of Fodor's central claim that it is impossible to learn a concept, I will ultimately conclude that we should be more sceptical than is normal about the power of learning when it comes to concept acquisition, particularly with respect to natural kind concepts. Central to my argument for this conclusion will be an examination of the bearings of psychological essentialism on concept acquisition.

Introduction .1

How do we acquire the concepts that we use to think? According to the most popular answer to this question in contemporary philosophy and cognitive science we learn most of our concepts. Such is the popularity of this answer that it is even endorsed by many who are generally quite nativist in orientation.¹ One persistent voice of dissent

¹ For example, Margolis (1998), Laurence and Margolis (2002), Margolis and Laurence (2011), Jackendoff (1989), Pinker (2007), and Carey (2009).

has been provided by Jerry Fodor. In this article I will evaluate the orthodox pro-learning view of concept acquisition through an examination of Fodor's reflections and those of some of his most persistent and significant critics. Although I will be critical of Fodor's central claim that it is impossible to learn a concept, I will ultimately conclude that we should be more sceptical than is normal about the power of learning when it comes to concept acquisition, particularly with respect to natural kind concepts.

Before turning directly to Fodor's views it is worth noting three points relating to the orthodox view. The first is that champions of the orthodox view of concept acquisition typically conceive of individuals as belonging to communities whose members share many concepts.² This in turn makes it possible for them to share a language on the assumption that what an individual means by a given word is a matter of what concept she uses that word to express. The second point is that champions of the orthodox view typically regard much of the environmental interaction involved in concept learning as a social phenomenon in the respect that it involves the learner interacting with other people.³ This interaction is based upon the would-be learner's desire to acquire the very concepts that her fellows use and the desire of her fellows to ensure that the would-be learner acquires the very concepts that they employ. The third point is a corollary of the first two: in many cases of concept learning it is not merely the case that a concept has been acquired on the basis of interaction with the outside world. In addition, the concept that is acquired is the very concept that the

² Jesse Prinz (2002) is clear example here as he argues that it is a condition on any good theory of concepts that it explains the fact that we share many concepts. He calls this the desideratum of publicity.

³ Sterelny (2003) constitutes a good example of this perspective.

learner had (however indirectly) intended to acquire or that that her fellows had attempted to direct her towards acquiring. Thus, there is an important respect in which concept learning is typically a reliable process: it reliably results in an individual acquiring the target concept, that is, the very concept she was aiming to acquire or her fellows were aiming for her to acquire. Consequently, learning ensures that members of a community share a common stock of concepts.

Fodor's classic argument against concept learning .2

In discussing concept acquisition Fodor's attention has focussed on lexical concepts, that is, concepts expressed by morphologically simple words. Following Cowie (1999) and Fodor (2001) we can distinguish between two types of argument against the view that we learn most of our lexical concepts. According to the former, it is impossible to learn a lexical concept so the very notion of such learning is incoherent. According to the latter, as a matter of empirical fact most of us never had experiences rich enough to facilitate the learning of our lexical concepts so we must have acquired them by some alternative means.⁴ In a number of publications Fodor (1975), (1981), (2001) has developed and defended an impossibility argument. The classic version of his argument can be expressed in the following manner:

P1) Learning a concept C is an inductive process of framing and confirming a hypothesis of the form 'x is a C if and only if x is an F'.

⁴ Such a view of concept acquisition echoes that of the acquisition of syntactic knowledge developed by Chomsky (1986).

P2) In order to frame such a hypothesis one needs the conceptual resources to
express the property F.

P3) If C is a complex concept then the property F can be expressed by using a
number of concepts each more basic than the concept C.⁵

P4) However, if C is primitive then one needs to employ the concept C in order to
express property F. The upshot of this is that with respect to primitive concepts one
needs to have a prior grasp of the target concept in order to frame the required
hypothesis.

C1) It is impossible to learn a primitive concept.

P5) The abject failure of philosophers and linguists to analyse any of our lexical
concepts in terms of more basic concepts suggest that most, if not all, of our lexical
concepts are primitive.

C2) Most, if not all, of our lexical concepts are not learned.

Those critics who have engaged with Fodor's argument have focussed their
attention on different elements of it. For example, Pinker (2007) rejects P5, Samet
and Flanagan (1989), Mandler (2004), (2008), and Margolis and Laurence (2011)
reject P1 and Margolis (1998), Laurence and Margolis (2002), and Weiskopf (2008)
reject P4.

⁵ For example, if C is the concept BACHELOR and being a bachelor can be analysed as being an
unmarried adult male, then one could frame the hypothesis needed to learn BACHELOR (namely, the
hypothesis that 'x is a BACHERLOR if and only if x is an unmarried adult male') without having a
prior grasp of BACHELOR.

I propose to evaluate the significance of Fodor's argument through an examination of Stephen Laurence and Eric Margolis's (2002) attempt to show that it is possible to learn a primitive concept. I will argue that although Laurence and Margolis establish that it is indeed possible to learn a primitive concept, there are reasons to be sceptical that the learning process that they postulate actually takes place in human cognitive development. I will then examine Fodor's (2008) return to the issue of concept acquisition in *The Language of Thought 2*. Here he argues that, contrary to his earlier view, the position with respect to complex concepts is no different from that with respect to primitive concepts. I will argue that there is an element of truth in this claim with the upshot that my scepticism with respect to the role of learning in the acquisition of primitive concepts carries through to the acquisition of complex concepts. However, there is a potential limit to the significance of these conclusions as the argument for them follows both Fodor and Laurence and Margolis in presupposing that concepts are symbols of the Language of Thought (LOT). This is potentially problematic as many cognitive scientists reject such a view of concepts. In the final stages of this article I address this worry by showing that my considerations apply just as much to pro-learning views of concept acquisition that endorse a non-LOT view of concepts.

Primitive concepts .3

Laurence and Margolis on primitive concepts 3.1

Margolis (1998) presents an account of how a primitive concept could be learned that has been developed and defended in a later article by Laurence and Margolis (2002)

and reiterated in Margolis and Laurence (2011). Laurence and Margolis focus on natural kind concepts for particulars (for example, DOG⁶). They argue that discussions of concept acquisition must take place within the framework of a commitment to a particular theory of content and the framework that they adopt is the atomistic theory of content developed by Fodor (1987), (1990).⁷ The account that they develop can be described in the following terms. Acquiring a concept involves acquiring a sustaining mechanism, that is, a mechanism that sustains the mind-world causal relations involving a mental representation that are such that that mental representation has the appropriate content. A kind syndrome for a particular kind 'is a collection of properties that is highly indicative of a kind yet accessible in perceptual encounters. This may include things like the typical shape, motions, markings, sounds, colors, etc., associated with a kind' (2002: 38). The sustaining mechanism for concepts such as DOG will include a representation of the kind syndrome. In other words, it will be a representational structure that carries information as to the perceivable properties that dogs typically have and so will be eminently learnable. Central to such a body of information will be shape information as dogs have a characteristic shape that distinguishes them from other types of animal. This body of information will be associated with a previously uninterpreted mental representation (a simple symbol of LOT for short) whose application it will drive. Thus, as the child gradually builds a body of information by means of learning the application of the representation moves closer and closer to meeting the conditions for having the content *dog* and thus being the child's concept DOG.

⁶ I am following the convention of using capitalised word to refer to concepts.

⁷ Though in the latter stages of their article they argue that their approach to concept acquisition could equally be developed in conjunction with a causal role theory like that championed by Block (1986).

The centrality of the representation of shape to the sustaining mechanism for DOG and other related kind concepts reflects an innate bias that is fundamental to the learning process. This is the so-called shape bias (Soja et al, 1991) that is such that when learning a concept primary attention is focussed on learning the typical shape of items falling under the target concept and the representation of shape, though it will not exhaust the representation of the kind syndrome, will be central to that representation and thus to the sustaining mechanism. However, as Laurence and Margolis continue:

shape by itself won't do, nor will any combination of simple perceptual features. Such features aren't a perfect guide to kind membership, since, among other things, the world is sometimes populated by what we'll call *fakes* – objects with the same outward appearance of a natural kind but which nonetheless aren't instances of the category. (2002: 40)

In addition, there is a second bias that constitutes a crucial element of the sustaining mechanism. This is the essentialist bias, the assumption made by the child that the items that fall under the concept are bound together by having a common essence. This essence is conceived as being hidden and being causally responsible for the perceivable properties that belong to the kind syndrome. The essentialist bias is a crucial part of the sustaining mechanism for kind concepts as without it the child would not acquire concepts such as DOG but more general concepts that apply to both dogs and 'fake dogs', that is things that look like dogs without being dogs. In sum, Laurence and Margolis think that they have presented a plausible account of how primitive kind concepts for particulars can be learned for:

Though the model requires a considerable amount of innate structure in the form of biases and inferential mechanisms of various sorts, [. . .] it still looks like a learning model. That's because it accounts for the acquisition of a concept which, in an important respect, reflects the character of one's experience. (2002: 43)

The view that children are essentialists (a view widely known as psychological essentialism) is prominent in contemporary developmental psychology being advocated by a range of scholars including Keil (1989), Gelman (2003), (2004), Bloom (2000), (2004), (2010), Pinker (1997) and Carey (2009). What this might be taken to suggest is that Laurence and Margolis are on firm empirical ground. Indeed, Susan Carey (2009) explicitly endorses their line of thought and states that the psychologist John Macnamara (1986) independently developed a very similar argument. However, I think that there are real problems with Laurence and Margolis's account of how we acquire primitive concepts. To say this is not to reject psychological essentialism but to say that their essentialist leanings do not help children to reliably learn concepts.

3.2 Evaluating Laurence and Margolis's argument

My critique of Laurence and Margolis will draw upon an examination of types of concepts that fall outside of those that they focus on, for example, kind concepts for stuffs. My justification for this is as follows. Laurence and Margolis's focus on kinds for particulars reflected the role of a shape bias that doesn't plausibly play a role in the acquisition of kind concepts for stuffs. However, there are potentially analogous

biases in relation to kind concepts for stuffs, namely, texture and colour biases. As Paul Bloom puts it: ‘different properties are relevant for the adaptive categorization of different entities. For rigid objects, for instance, shape is highly relevant. . . . But for substances color and texture are what matter’ (2000: 149). Given this, along with the fact that many advocates of psychological essentialism think that children are essentialist about stuffs, it is difficult to see why Laurence and Margolis’s account would be viable with respect to one type of kind concept but not the other.

Following Medin and Ortony (1989) advocates of psychological essentialism typically argue that children have a placeholder conception of essence; that is, they don’t have a concrete view of the nature of the essence in question. But could a commitment to essentialism play a central role in concept learning if it was so unspecific? Both Gelman (2003) and Carey (2009) draw an explicit comparison between psychological essentialism and externalist views of meaning and reference developed by Putnam (1975) and Kripke (1980) so it will be helpful to consider the issue through a discussion of Putnam’s account of how meaning is bestowed upon words such as ‘water’.

According to Putnam the word ‘water’ got its meaning through a baptism ceremony where an individual pointed at a sample of water and thought ‘I will call that kind of stuff “water”’. In other words, she thought of the sample of liquid she was pointing to as belonging to a natural kind and wanted to attribute to ‘water’ a meaning such that it applies to, and only to, instances of that kind. But she didn’t think anything more specific than this; for example, she didn’t think that instances of the kind are bound together by having a common microphysical essence. Putnam thinks that this was enough for the naming ceremony to succeed in bestowing the meaning *water* on ‘water’, a meaning such that that word applies to, and only to,

samples of H₂O. I'm not so sure. My worry here is that because the individual's state of mind was so unspecific she will not have pinned down what counts as 'being the same kind of stuff' as the sample of water she points at. And as a consequence of this she will not have bestowed a determinate meaning on the word 'water'.

It might be objected that it's a metaphysical fact that the stuff that the individual interacted with had an essence residing at the microphysical level and that that was enough to ensure that she attributed the meaning *water* to 'water'. However, there is a potential problem when the individual is interacting with an item that falls under two distinct kinds with quite different forms of essence. Consider a sample of milk. The sample of milk in question will have a particular microphysical structure and chemistry; it will largely be H₂O but will also contain various fats, proteins, minerals, vitamins, and the like. But its essence qua milk cannot be identified with such microphysical and chemical properties; something could be microphysically and chemically quite unlike this sample of milk whilst being milk and something could be microphysically and chemically identical to the sample without being milk. To see this consider the following thought experiment. Suppose that a distant planet is inhabited by a species of super-intelligent reptiles and has never been visited by mammals. In their laboratories the reptiles synthesise a liquid that is microphysically and chemically just like cows' milk. Would this stuff be milk? No, for the simple reason that it wasn't produced naturally within the body of a mammal-like creature for the purpose of feeding and sustaining its young. Now consider Twin Earth. The twin-cows there produce a liquid in their body that they use to feed their young that is quite unlike cows milk at the microphysico-chemical level as it is largely made up of XYZ. Yet surely this stuff is milk. What this shows is that milk has a quite different kind of essence than water; one having to do with origins and function rather than

straightforward microphysical structure and chemistry. But any sample of milk will also be a sample of a liquid with particular microphysical and chemical properties. The upshot of this is that if an individual in a Putnamian naming ceremony points at a sample of milk and says ‘I’ll call that kind of stuff “milk”’ they haven’t thereby succeeded in bestowing the meaning *milk* on ‘milk’. This is because they haven’t pinned down what it is to be ‘that kind of stuff’. Is being ‘that kind of stuff’ to do with microphysical structure and chemistry or is it to do with origins and function? Unless the individual has something in her mind which settles this question then she will not succeed in bestowing a determinate meaning on ‘milk’ and if she does have such an answer in her mind then the consequences of the naming ceremony will depend on the specific form that that answer takes.

Examples with the same moral exist within the realm of words for particulars. To be a ‘dog’ in this sense is largely to do with being an animal that belongs to a lineage that has been selectively bred by humans from wolves to perform one of a range of functions (for example, guarding, herding, retrieving, hunting, and so on) within a domestic setting. Consequently, DOG isn’t a purely biological concept. Rather, it is a quasi-artifactual concept. Now suppose that an individual in a Putnamian naming ceremony encounters a domestic dog and points at it saying ‘I’ll call that kind of animal “dog”’. What meaning they will have bestowed upon ‘dog’ will depend upon the notion of essence they have in mind. If they think of ‘being that kind of animal’ as being a matter of belonging to the same biological kind as the animal before them then the meaning *canine* will be bestowed on ‘dog’ so that that term will apply just as much to foxes and wolves as it will to domestic dogs. However, if they think of the essence in quasi-artifactual terms, that is, as having to do with the history and role in human affairs then a different meaning will be bestowed upon ‘dog’. And if the

individual has no determinate notion of the essence in mind then no determinate meaning will be bestowed upon 'dog'.

I will now bring these points to bear on Laurence and Margolis's account of concept learning. If a child is going to bring a commitment to essentialism to bear in acquiring a range of concepts then she needs more than a single undifferentiated notion of essence. Rather, she will need to have a range of worked out notions of essence and make the appropriate selection of which such notion to bring to bear in each particular case. This is because essences of the categories for which we have concepts come in different forms. Some categories have microphysical essences, some have biological essences and some have artefactual essences, to identify three forms of essence. And some essences are a mixture of these things. It is not just that essences come in different forms but that individual items (things or samples of stuff) can fall under distinct concepts each of which picks out a category with a distinct form of essence.

This raises two related problems for Laurence and Margolis. First, is it really plausible that children have a range of distinct articulated notions of essence? For wouldn't such notions of essence require them to have a grasp of a range of technical concepts relating to microphysical, biological and artefactual properties? And if so, how were those concepts acquired? Second, and more importantly, even if we were happy with the idea that children had such a range of notions of essence that they were able to employ at the age when they were acquiring concepts such as MILK and DOG what evidence does the child have as to which notion of essence to bring to bear in each particular case? What I am raising here is a poverty of the stimulus worry that can be developed in the following terms.

I have argued that MILK picks out a category that doesn't have a microphysical essence and thus contrasts with WATER. We might say that the essence of milk is at least partly biological. If a young child is to learn the concept MILK by utilising an essentialist commitment then she needs to know that the category referred to in her community by the word 'milk' has a biological essence rather than a microphysical essence. And if she is to learn the concept WATER by utilising an essentialist bias then she needs to know that the category referred to in her community by 'water' has a microphysical essence. Any mistakes here will result in her acquiring the 'wrong' concept, that is, a concept at odds with that expressed by 'water' on the lips of most of her fellows. But what tells the child what the relevant form of essence is? Presumably children are rarely told that milk has a biological essence whereas water has a microphysical essence. For, as Paul Bloom points out:

Even highly educated parents in university towns rarely talk to their children about insides and essences . . . and working-class parents are considerably less likely to do so. (2000: 154)

Could children find out about essences less directly? I do not wish to deny that children receive lots of information about both 'milk' and 'water' when they are acquiring the corresponding concepts. But the problem that the child faces is in determining the significance of that information. A child may well receive plenty of information about the typical origins and function of the stuff her fellows call 'milk'. For example: that 'milk' is made in the bodies of animals such as cows and humans; that it is the primary food source of the young of such animals; that much of the 'milk' that we drink after early childhood is taken from the bodies of cows and more

immediately acquired from shops and supermarkets where it is packaged in cartons and bottles, and so on. But what tells the child which of these features of 'milk' are central to its essence and which not? What tells her that (biological) origins and function are central to essence but that microstructure isn't so crucial? Similarly, the child will receive plenty of information about what her fellows call 'water' some of it relating to its origins and its use by biological systems. For example: 'water' falls as rain and fills rivers, lakes and seas but is made more immediately available to us through taps; that 'water' is fundamental to life in that animals and plants need to consume 'water' to survive; and so on. Much of the information about 'water' and 'milk' that children receive is analogous in that it relates to the respective role of these stuffs in human and animal life. My point is that this raises the question as to how children could reach different conclusions about the relevant notion of essence involved in each case. Why would a child think 'water' has an essence that does not relate to its role and function in animal life when much of what she learns about 'water' is that it has a crucial role and function in animal life? And why would she think that the role and function of 'milk' in the lives of those animals that consumed it was central to its essence? My point is that the information freely available to children could just as well lead them to think that 'water' and 'milk' had comparable essences as that they had divergent types of essences. Therefore, Laurence and Margolis haven't presented a plausible account of a mechanism that enables children to reliably acquire the 'correct' concepts on the basis of their experiences, that is, the concepts possessed by the adult members of the community and expressed by the words 'water', 'milk', and so on.

This is not to say that children could not acquire determinate concepts on the basis of their experiences of milk and water. However, there would be nothing to ensure

that the concepts that one child acquires agrees with those of the next or with those of the adults whose concepts she is attempting to acquire. Similar considerations apply to concepts such as DOG. Recall that I argued that DOG picks out a category that doesn't have a purely biological essence, that DOG is a quasi-artefactual concept. But children acquire plenty of other concepts for animals that are not quasi-artefactual. The concept RAT is an example. Now consider what a child learns about the animals referred to by 'rat' as she is acquiring the concept RAT. She learns a lot about the lifestyle of 'rats', how their history has been bound up with that of humans and of their role in human life. But in the case of RAT she had better not think that these properties to do with the role and history of 'rats' in human life, though highly significant, relate to the essence of 'rats'. For if she does she will not acquire the concept RAT but some distinct quasi-artefactual concept. But in the case of DOG she must think that properties relating to the history and role of 'dogs' in human life are central to their essence otherwise she won't acquire the concept DOG but some purely biological concept. In short, the child would need to bring to bear distinct notions of essence when learning these distinct concepts and we should be sceptical that children typically receive information that reliably indicates which notion of essence to utilise in the case of which concept.

What I have argued is that Laurence and Margolis may well be right in arguing against Fodor that there are certain primitive concepts that it is in principle possible to learn by utilising an essentialist bias. But on the basis of considering both concepts for stuffs such as WATER and MILK and concepts for particulars such as DOG and RAT, I have argued that when we consider their account of learning in more detail poverty of the stimulus considerations emerge. That is, there are reasons for believing that children don't generally have rich and specific enough experiences to reliably

learn concepts such as MILK, WATER, DOG and RAT by utilising a commitment to essentialism.

A first response that might be made on behalf of Laurence and Margolis is that all I have done is shown that the mechanisms they describe can't support the reliable learning of concepts; that is, that they bring an element of randomness into the concept learning process. But, so the objection continues, learning doesn't require reliability so that, for example, an individual could have learned the concept that she expresses by the word 'water' even if similar experiences would have lead some of her fellows to acquire a different concept. In reply to this objection I would emphasise the point made in section one that it is part of the orthodox view of the role of learning in concept acquisition that learning is, as a matter of fact, a reliable process that ensures that members of a community share a common stock of concepts that they express by means of a shared language. The upshot of this is that my argument implies that Laurence and Margolis's theory of concept learning cannot be employed in the service of the orthodox view.

It should also be pointed out that Margolis and Laurence (2011) commit themselves to the view that the learning mechanism they postulate is reliable in responding to an objection delivered by Fodor. Fodor (2008) argues that even if learning a sustaining mechanism is sufficient for acquiring a concept it doesn't follow that the concept is learned. His worry is that there are cases where the sustaining mechanism that locks the acquired concept onto the property it expresses is irrational, incoherent or untrue. He presents the Ancient Greeks' theory that stars are holes in the heavenly firmament, a theory that served to lock their concept STAR onto the property of being a star, as an example of such a sustaining mechanism. In such cases, thinks Fodor, it is an accident or fluke that learning the sustaining mechanism

facilitates the acquisition of the concept in question and for that reason the acquired concept isn't learned. Margolis and Laurence's (2011) rejoinder is that the existence of cases such as those Fodor describes does not establish that the problematic accidental element is present in all cases where the construction of a sustaining mechanism leads to the acquisition of a concept. For, acquiring a sustaining mechanism could involve gathering correct information by means of a mechanism that has the function of gathering such information and reliably performs that task. Such a case would count as learning a concept even if the Greeks didn't learn their concept STAR. The relevant point here is that Margolis and Laurence claim that the concept acquisition mechanism that they postulate is typically reliable as opposed to random or accidental.

A second response that might be made on behalf of Laurence and Margolis is that their key point is that it is possible to learn a natural kind concept and that this doesn't require them to provide an account of how we actually learn such concepts. Consequently, my argument misses its target by being based upon a misunderstanding of their general goal. My response to this charge is that the issue of how we acquire concepts both is and should be important to Laurence and Margolis. As we have seen, they have done valuable work examining Fodor's view that it is impossible to learn a concept. Part of their explicit motivation for carrying out such work against a general background of dismissive hostility to Fodor's views in this area is that they think that a detailed engagement with Fodor helps to shed light on the nature of the role that learning plays in the human concept acquisition process. It is inconsistent with this approach to be unconcerned with the specifics of how we actually acquire concepts. Moreover, surely the important issue is how we acquire concepts and merely establishing that their acquisition by means of learning is in principle possible only

sheds limited light on this issue. For, the mere in-principle-possibility of concept learning doesn't imply that we actually learn concepts. Ultimately then, the real interest in Laurence and Margolis's work relates to the question of whether they have given us good reason to think that we humans actually learn our (natural kind) concepts.

I now come to the issue of the relationship of my position to that of Fodor as it might be argued⁸ that there is little to distinguish my argument from his and that he would welcome my invocation of empirical considerations to bolster an anti-learning perspective. For, the charge continues, my point is in effect that, for example, if a child does not employ the concept WATER when attempting to learn that concept then she will acquire a concept with an indeterminate content. This accusation is mistaken. The key difference between my position and Fodor's is that I accept that it is possible to learn a natural kind concept such as WATER, DOG or MILK and so I do not endorse Fodor's impossibility argument. In arguing against Laurence and Margolis I claim that a child must bring more to her interaction with instances of the concepts she ultimately acquires than a generic essentialist bias. But that more specific something is not the target concept. For example, in learning the concept WATER by utilising a commitment to essentialism a child would need to employ the concept of a microphysical essence. Hence, if a hypothesis was being framed in this process it would take something like the following form:

A sample of liquid is WATER if and only if it has the same microphysical constitution as the local liquid that falls as rain, comes out of taps, and so on.

⁸ As it was by an anonymous referee for this journal.

The key concept in this hypothesis is MICROPHYSICAL STRUCTURE. This concept is clearly not to be identified with that of WATER and can in principle be used in learning a series of distinct natural kind concepts, namely those whose instances are bound together by sharing a common microphysical structure.

Complex concepts .4

The *LOT2* development 4.1

I now turn to the issue of complex concepts. Recall, that in his classic argument for concept nativism Fodor conceded that complex concepts can be learned but argued that most lexical concepts are in fact primitive rather than complex. However, in *LOT2* Fodor modifies his position by rejecting the idea that complex concepts can be learned with the upshot that his anti-learning argument no longer relies on the premise that most lexical concepts are primitive.⁹ Hence, even if lexical concepts were analysable in terms of simpler concepts that wouldn't imply that they were learned. Within the context of the discussion in this article, this change of heart is potentially very significant. For, if Fodor's reasoning is correct then my scepticism with respect the role of learning in the acquisition of primitive concepts will imply that we should be equally sceptical about the role of learning in the acquisition of complex concepts. For this reason I now turn to Fodor's new claim that it impossible to learn a complex concept.

⁹ It is this premise that has come under substantial attack, particularly from linguists working in lexical semantics (Pinker, 2007; Jackendoff, 1989).

Fodor's argument for the view that there is no relevant difference between the acquisition of primitive and complex concepts is straightforward. Suppose that the target complex concept C is the concept GREEN OR TRIANGULAR. To learn this concept one has to frame and confirm a hypothesis of the form 'x is a C if and only if it is green or triangular'. But framing and confirming this hypothesis involves using the concept GREEN OR TRIANGULAR and this is clearly something that one cannot do unless one already has that concept. Therefore, it is impossible to learn a complex concept.

Is this development in Fodor's position to be welcomed? In the latest of their important series of articles criticising Fodor's views on concept acquisition Margolis and Laurence (2011) argue for a negative answer to this question and it will be instructive to examine their line of thought.

Laurence and Margolis's counterexample 4.2

Margolis and Laurence attempt to develop a counterexample, a case of a complex concept that is clearly learned. An intermediate student of Scottish Country Dancing is attempting to learn what is involved in a particular dance known as 'Maxwell's Rant' on the basis of watching dancers perform this dance. The dance is complex consisting of eight particular moves arranged in a specific order.¹⁰ The student is

¹⁰ More specifically, Maxwell's Rant is:

a dance that involves the following sequence: *reflection reels of three on opposite side, followed by reflection reels of three on own side, followed by crossing with right hands, followed by casting off, followed by half figure of eight, followed by leading down the set,*

familiar with all these moves and has concepts for each of them. Her task is to build a complex representation of the constituent moves of Maxwell's Rant and their ordering. Completion of this task would constitute acquiring the (complex) concept MAXWELL'S RANT. The question is: can this be done by means of hypothesis testing? Margolis and Laurence argue for an affirmative answer. They describe the point in this process where the learner has worked out seven of the eight moves and their ordering but is yet to determine the penultimate move. Watching another performance of the dance she finally identifies the penultimate step and at that point appreciates that she is now in a position to construct a complete and accurate representation of the dance. For Margolis and Laurence, what is crucial is that at this point the dancer has done all the hypothesis testing (all the gathering and reflecting on evidence that she needs) but is yet to construct an explicit representation of the target concept. Therefore, that final step in the process of acquiring the concept comes after the hypothesis testing has taken place so that executing the relevant hypothesis testing process does not presuppose already having the target concept. Thus, in this case, the hypothesis testing model of concept acquisition that Fodor criticises is not circular. Margolis and Laurence contend that such cases where the target complex concept is explicitly represented only after all the hypothesis testing has taken place are ubiquitous and involved in cases where the target concept is a concept 'that describes a sequence of events in terms of a more basic stock of event types (e.g. concepts involved in learning a chess strategy, a cooking recipe, a new type of knot, or a chord change)' (2011: 513).

followed by casting up, followed by turning with right hands (Margolis and Laurence, 2011: 511).

I am not convinced that this is a genuine counterexample to Fodor for two reasons.

First, exactly when did the dancer acquire the concept MAXWELL'S RANT? For Margolis and Laurence to have presented a genuine counterexample it has to be the case that the target concept was acquired only at the point at which the explicit complex representation was formed. But one might argue that the concept has been acquired well before then. Consider the point when the dancer identified the previously unknown penultimate move. At this point she has two distinct representations in her mind. One represents the dance as consisting of eight components and represents the identity of seven of those minus the penultimate move. And the other represents the penultimate move. In virtue of the dancer's readiness and ability to integrate these representations to form a single complex representation of the nature of the dance I would argue that she has the concept at this point prior to executing any such integration.

My second reason for doubting that Margolis and Laurence have produced a genuine counterexample relates to the worry that the dancer possessed the concept MAXWELL'S RANT prior to her beginning the process of explicitly constructing the complex representation that they describe. Thus, what she acquires as a result of observing the performances of her fellow dancers is not a new concept but something else; for example, knowledge that the particular dance she has been observing is called 'Maxwell's Rant' or that that dance is a particularly important or oft-performed routine in the Scottish Country Dancing community.

In 'The present status of the innateness controversy' Fodor (1981) commits himself to the view that both primitive concepts and the mechanisms (rules) for

combining them to make complex concepts are innate and so not learned. This innate endowment sets the limits of the concepts we are capable of thinking. Yet he goes on to say that it is only in a trivial sense that complex concepts are innate. The idea seems to be that when we explicitly token such complex concepts that counts as a case of having acquired them (typically by means of learning) and the point of acquisition is the point in time when the complex was first explicitly tokened in the mind-brain. Margolis and Laurence seem to endorse this point as is indicated by the following quote:

someone who encounters a black swan for the first time is likely to form the concept BLACK SWAN. She needn't have had the concept prior to her encounter, and it might never have occurred to her that non-white swans are a real possibility. But if she has the concepts BLACK and SWAN and sees a black swan with her own eyes, she will come into the possession of BLACK SWAN (2011: 519).

On this point I disagree with both Fodor and Margolis and Laurence: there are plenty of cases where an individual possesses (and so has already acquired) a complex concept prior to her first explicitly tokening that complex. To see this consider the following. Suppose that prior to time t , though I have often thought of aardvarks and often thought of particular animals as being angry, I have never thought of anything as being an angry aardvark. Thus, I have never explicitly tokened the concept ANGRY AARDVARK prior to time t . At time t time I am watching a nature programme about aardvarks and witness some overtly grumpy behaviour and come to think 'THAT'S ONE ANGRY AARDVARK' so explicitly tokening the concept ANGRY AARDVARK for the first time.

In this context it is very important not to lose sight of Fodor and Margolis and Laurence's commitment to the Language of Thought hypothesis and with it the idea that concepts are symbols of the LOT (words of LOT in the case of primitive concepts and phrases of LOT in the case of complex concepts). For a champion of LOT there is a strong analogy between thought and language as viewed within the Chomskyan generative tradition in linguistics. One of the key phenomena that Chomsky has been keen to explain is the so-called creativity of language.¹¹ That is, the fact that we routinely produce and hear sentences that we have never encountered before but have no problems understanding these sentences and appreciating their grammaticality. This is explained by postulating a state of knowledge of language (an I-Language) that comprises of a lexicon and a battery of recursive rules or principles for combining such lexical items and moving them around within complex structures.¹² Now suppose that an individual hears a particular sentence for the first time (for example, 'Hedgehogs carry too many fleas to make good household pets') and finds no trouble understanding it. This is a consequence of her knowledge of the relevant lexical items and syntactic rules of English. But it would very odd to say that the when the individual first hears and understands this sentence her language or her knowledge of the language that she speaks has expanded. She has not added a new sentence to her language or acquired the knowledge that the sentence in question belongs to her language. There has been no linguistic development in this case and therefore no learning. By parity of reasoning, assuming LOT, when an individual explicitly tokens a complex concept for the first time (for example ANGRY

¹¹ See, for example, Chomsky (2009).

¹² See Boeckx (2010) and Isac and Reiss (2013) for an accessible account of this Chomskyan perspective on language.

AARDVARK or BLACK SWAN) by drawing upon the possession of the constituent concepts of that complex and (her knowledge of) rules for combining simple concepts she doesn't add anything to her conceptual scheme. Therefore, there has been no development with respect to her conceptual scheme and so no learning. To reject this claim is to lose sight of the analogy between language and thought within the LOT perspective.

Let's now return to Margolis and Laurence's example. The dancer acquires the concept MAXWELL'S RANT during the period of observing the dance being performed. The putative acquisition of this concept involves the construction of a complex representation that represents the elements of the dance and their ordering. The elements of the dance are moves for which the dancer already has concepts and in combining those concepts in the manner she does she employs combining mechanisms that she has employed many times before. This suggests that the case is just like that of ANGRY AARDVARK so that the dancer does not acquire the concept MAXWELL'S RANT at the point at which she first explicitly constructs the complex representation Margolis and Laurence describe. On the contrary, she had the concept all along. If this is correct then Margolis and Laurence do not have a counterexample to Fodor at all.

I can envisage a number of objections to this line of reasoning. The first is that the case of ANGRY AARDVARK is not analogous to that of MAXWELL'S RANT as in the latter case the explicit tokening of the concept did not come as easily as in the former. All I had to do was see film footage of an angry aardvark to be prompted to explicitly token ANGRY AARDVARK. In contrast, the dancer had to carefully watch several performances of the dance and gradually build up the relevant complex representation. I don't think we should be impressed by this difference. Throughout

the putative learning process the dancer knows that MAXWELL'S RANT is a dance consisting of familiar dance moves placed in a particular order. The problem she faces is in identifying the particular moves involved in the dances she witnesses. Thus, her problem relates to applying concepts that Margolis and Laurence accept that she has had all along. This is presumably because the dancers are moving swiftly and each stage of the dance quickly merges into the next. Perhaps she wouldn't face such a problem if she watched a slow motion film of the dance.

A second objection appeals to long-term memory and runs as follows. Once the complex concept that Margolis and Laurence describe has been explicitly constructed it can be stored in long-term memory. This facilitates the future application of the concept, aids its lexicalisation, and provides a locus for building a body of detailed information about the dance. This transition is significant and motivates the idea that for a complex concept to be genuinely possessed it has to have been explicitly constructed and stored in long-term memory. In response to this objection I accept that the transition described is significant but reject that it counts as a case of concept acquisition or learning. What has been acquired or learned is not a concept but something else; for example, a commitment to the view that a particular concept is significant and therefore likely to be expressed by a word, to figure in distinctive generalisations worth knowing, and such like. In short, we have to be clear on what is learned and it is not always a concept in the cases that Margolis and Laurence consider.

A further response to the objection involves pointing out that within the LOT model of concepts no inevitable commitment to the idea that when an individual has full possession of a concept there is a corresponding symbol of LOT explicitly stored in her long-term memory. This can be seen by considering Fodor's view of the

primitive symbols of LOT that belong to an individual's innate endowment. Suppose that this includes the concept DOG. Then there is a primitive symbol of LOT that is locked onto the property of being a dog. What this requires, to a first approximation, is for there to be a primitive symbol of LOT such that its tokening is caused by and only by dogs.¹³ For this to be the case, such a symbol only needs to appear in the mind briefly whenever the individual is impinged upon by a dog. This no more requires the symbol to be stored in long-term memory than my tendency to experience pain when punched on the nose requires a pain to be stored in my long-term memory. But if one accepts that one can fully possess a primitive concept without having a corresponding representation explicitly stored in long-term memory then it is difficult to see how one could legitimately insist that a complex concept is fully possessed only when it is explicitly stored in long-term memory. Recalling the point made in the previous paragraph, to say this is not to deny that complex concepts are often explicitly stored in long-term memory or that there can be cognitive value in doing this. Rather, the point is that such storage isn't constitutive of possessing the concept but reflects the individual's commitment to the potential significance of the concept in question.

To argue that complex concepts are possessed in virtue of the possession of their simple components and combination rules prior to their having been explicitly tokened is not to imply that, within the LOT framework, the same can be said of all representational states. For example, consider beliefs. Suppose someone has never explicitly considered the proposition that hedgehogs have too many fleas to make good household pets and so has never tokened the LOT sentence HEDGEHOGS

¹³ This view of the semantics of mental representations was developed and defended by Fodor (1987), (1990) drawing upon earlier work by Dretske (1981).

HAVE TOO MANY FLEAS TO MAKE GOOD HOUSEHOLD PETS. Nevertheless, in virtue of the nature of their LOT they have the capacity to token this sentence in forming a belief or some other type of propositional attitude. But it doesn't follow from this that they actually possess the propositional attitude in question. In other words, in the case of propositional attitudes one has to explicitly token a relevant representation in order to have that attitude. The mere possession of the resources that would enable one to token the representation if suitably prompted is not enough. This difference between complex concepts and propositional attitudes is solidly grounded within the LOT framework. Propositional attitudes have causal powers; they have a causal impact on an individual's behaviour and her mental states. What explains this is that when one has a propositional attitude one tokens a sentence of LOT in a relevant functionally individuated component of one's mind-brain (for example, in one's belief box). But concepts in themselves don't have causal powers quite the same way. It is not my possession of the concept DOG that causes me to token certain mental states or behave in a particular way. Rather, it is particular applications of that concept in categorisation or thought that have such effects, applications that do involve the explicit tokening of the relevant symbol of LOT.

My discussion of Margolis and Laurence's putative counterexample has bearings on another important element of their critique of Fodor. This is their objection to his claim that concept learning always involves hypothesis testing. Following Fodor (1981), (1998) they accept that learning is a rational-causal rather than a brute-causal process but argue that a rational-causal process of concept acquisition need not involve hypothesis testing. Fodor's mistake, they contend, is to endorse an internalist theory of justification. Once one countenances an externalist theory of justification then it becomes clear that there are three familiar means of acquiring a concept that

are suitably rational-causal to count as learning where the learner does not frame a hypothesis and then consider evidence for its truth. The first of these is perceptual learning as when:

someone forms a new complex concept as a result of perceiving an object or event that manifestly exhibits the combination of properties that the concept picks out. For instance, someone who encounters a black swan for the first time is likely to form the concept BLACK SWAN. She needn't have had the concept prior to her encounter, and it might never have occurred to her that non-white swans are a real possibility. But if she has the concepts BLACK and SWAN and sees a black with her own eyes, she will come into the possession of BLACK SWAN and will be prepared to record and organize new information about these unexpected creatures. (p. 519)

As should be clear from the foregoing, I would not accept that this is a case where a new complex concept is acquired let alone learned. Rather, the individual's ability to readily identify the creature before her as a black swan is a product of her prior grasp of the concept BLACK SWAN. Insofar as she learns anything it is not a concept but the information that there are black swans (that something actually falls in the extension of the concept BLACK SWAN) or that the category of black swans is a potentially significant one.

The second means of learning a concept that Margolis and Laurence identify is communication-based learning where a concept is acquired by means of 'verbal instruction and learning' (p. 519) as when the concept of a VALID ARGUMENT is explained to a student. One might wonder how much the existence of this kind of

learning would worry Fodor as in *The Language of Thought* (Fodor, 1975) he noted that such learning does take place and that it doesn't fit the hypothesis testing model.

However, he seems to regard it as implausible that any more than a small number of our concepts are acquired in this way and on this matter he is surely correct.

Moreover, it is arguable that the case where the concept VALID ARGUMENT is explained to a student is not one where she thereby acquires or learns a new concept.

Rather, what she learns is what the expression 'valid argument' means and her prior possession of the concept VALID ARGUMENT was indicated by her ability to

reliably distinguish valid from invalid arguments.¹⁴ In other words, once again, we mustn't lose sight of the distinction between acquiring a concept and learning which

concept is expressed by a particular word or phrase of a public language.

The third kind of learning is what Margolis and Laurence call 'automatic

associative learning' as when a cognitive system:

separately monitor[s] the environment for pairs of properties, F and G, and . . . [is]

. . . designed so that if it registers that a significant number of objects have both of

these properties while few objects have only one of these properties on its own,

then the system will form a new complex concept that encodes their conjunction (F

& G). (p. 520)

The problem here is that this no more looks like a case of acquiring a new concept

than the case of BLACK SWAN for precisely the same reasons.

¹⁴ In my experience the successful teaching of logic involves getting students to rigorously and consistently apply and represent reasoning skills they already have.

In sum then, Margolis and Laurence have hardly given us given us grounds to reject the Fodorian view that learning plays little or no role in the acquisition of complex concepts.

5. Alternative theories of concepts

5.1 Prototypes and proxytypes

The argument that I have developed in the previous sections follows both Fodor and Laurence and Margolis in presupposing the view that concepts are symbols of LOT. But, it might be objected, this limits its significance as many cognitive scientists endorse alternative views of concepts. In this section I address this objection by examining the application of my considerations to alternative views of concepts. A first alternative is presented by the prototype theory. Advocates of the prototype theory regard concepts, in the form of prototypes, as mechanisms of classification; that is, they are complex representational structures that are the means by which we categorise objects that impinge upon us.¹⁵ A consequence of this is that concepts need to be explicitly tokened in the mind-brain (in the form of some physically realised state) if they are to do their constitutive job. Therefore, if one merely has the potential to construct a particular prototype in one's mind-brain one does not thereby possess the corresponding concept. For Fodor, in contrast, concepts are not mechanisms for classification; rather, they are the representations that are explicitly tokened when such mechanisms are employed. It is because of this that one can have a concept as part of one's conceptual repertoire whilst one is not explicitly tokening that concept (just as one can have a sentence as part of one's language whilst one is not speaking

¹⁵ This feature of the prototype theory is emphasised by both Prinz (2005) and Machery (2009).

or hearing it). The upshot of this is that my argument that complex concepts are possessed prior to their being explicitly tokened won't work if concepts are prototypes.

Moreover, if we examine the process by means of which prototypes are constructed I think it is clear that they are learned. To see this consider the following.

The construction of a prototype is a process extended in time where a complex representation is gradually constructed by adding representations already belonging to the individual's representational repertoire to a provisional representational structure and/ or adjusting the weightings of such elements.¹⁶ The end result is a representation

that corresponds to a category that no simpler representation available to the individual at the beginning of the construction corresponds to. Consider, for example,

the prototype corresponding to the concept DOG. It will be constructed from representations that represent properties that dogs often have, such as those of being a quadruped, of barking, of having a tail, of being covered in fur, and so on. The

selection of representations of these properties, as opposed to representations of properties that dogs never or rarely have, along with the assignment of their

weighting, will be determined on the basis of relevant experiences, for example, those involving the perception of dogs. Echoing Margolis and Laurence, I would say that

this process counts as one of learning as the construction process is sensitive to relevant features of the environment and gradually constructs a prototype that

facilitates the reliable categorisation of dogs. It is also far from clear that it is one where hypotheses are framed and tested. At each stage in the construction process a

complex representation is modified on the basis of evidence as to its adequacy and by this means the target prototype is eventually arrived at. But it is not as if the

¹⁶ See Murphy (2002) for a detailed account of such a process.

representation constructed at each stage is floated as a tentative hypothesis with the explicit intention of testing it against further data that is actively gathered for that end.

Rather, it is just that the representation is not fossilised so that it can be modified in the light of subsequent experiences should those experiences indicate that it is inadequate. The upshot of all this is prototypes can be learned and that their learning need not involve hypothesis testing.

Extending the argument to protoypes and proxytypes 5.2

Fodor is strongly opposed to the identification of concepts with prototypes as he argues that prototypes, unlike concepts, do not generally compose.¹⁷ I do not wish to pursue this line of objection. Rather, my objection to the prototype theory is that it has

difficulties handling natural kind concepts. These difficulties emerge when we examine a close relative, namely Jesse Prinz's (2002), (2005) proxytype theory.

For Prinz we do indeed learn concepts such as WATER. Doing so involves constructing a complex network of representations on the basis of experiences of water. This network is stored in long-term memory and is ultimately made up of perceptual representations. Whenever an individual employs the concept WATER an element of this network is activated in short term memory and so goes proxy for the whole network from which it is drawn. Hence, it is a proxytype. That such a proxytype and the larger network from which it is drawn has the content *water* is due to the fact that the latter was constructed on the basis of perceptual interactions with samples of water.

¹⁷ See Fodor (1998) and Fodor and LePore (2002).

As both Carey (2009) and Bloom (2000) point out, the problem for the standard prototype theory of concepts is that they don't do justice to the nature of the concepts that we actually have. Children just don't think that falling under the concept DOG is merely a matter of having a statistically significant proportion of the perceivable properties that dogs have had in their experience. Rather, they think that falling under the concept DOG is something that goes beyond or behind appearance. And it is this that suggests that any attempt to explain how we learn those concepts that we acquire in early childhood is going to face a major challenge from the kind of considerations that I have directed at Laurence and Margolis. To see this let's return to Prinz's proxytype theory.

Putnam's (1975) Twin Earth thought experiment generates a problem for Prinz, namely that of explaining how Oscar on Earth and Twin-Oscar on Twin Earth could have learned divergent concepts. More specifically, how could the concept that Oscar comes to express by means of 'water' diverge in content from that that Twin-Oscar expresses by the same word? The problem for Prinz is that given that the representational structures that are stored in long-term memory from which proxytypes are drawn are ultimately constructed out of perceptual representations, it would appear that the twins have exactly the same proxotypes and, therefore, exactly the same concepts.

Prinz is alive to this problem and in addressing it he employs Locke's distinction between real and nominal essences. The real essence of water (that is, the colourless, odourless liquid found here on Earth) is a matter of its microphysical constitution. The nominal essence of water is a matter of the perceivable properties characteristic of water on the basis of which we typically identify a sample of water as such. Corresponding to this distinction is that between real and nominal content. The real

content of the respective concepts expressed by means of 'water' by Oscar and Twin Oscar differ. This is because the stuff falling under Oscar's concept has the real essence of being H₂O whilst the stuff falling under Twin Oscar's concept has the real essence of being XYZ. On the other hand, their concepts have the same nominal content as the perceptual representations that figure in the proxytypes that constitute their respective concepts are identical.¹⁸ In effect, what Prinz is saying is that the real content of a particular concept possessed by an individual is a matter of the essence of the items that the individual causally interacted with in constructing that concept. As Oscar interacted with H₂O in constructing his concept, that concept has the real content *water*. Whereas, Twin Oscar's corresponding concept has the real content *twin water* as it was constructed on the basis of casual interactions with Twin Water.

However, what I have said so far leaves out a crucial aspect of Prinz's line of thought and this is his endorsement of psychological essentialism. Thus, with respect to Oscar Prinz would say that he thinks of the stuff falling under his concept WATER as having a particular essence (the nature of which he may well think himself ignorant) that is the causal basis of the perceivable properties in virtue of which he typically identifies a sample of water as such (that is, the properties that are represented by the relevant proxytype). Thus, Prinz accounts for the real content of Oscar (and our) concept WATER on the basis of Oscar's (and our) essentialist commitments along with the fact that that concept was constructed on the basis of causal interactions with H₂O. Without such an essentialist commitment the concept Oscar and we express by means of 'water' would have a content such as to apply to

¹⁸ Prinz says that his distinction between real and nominal content corresponds to the familiar one between broad and narrow content.

anything with an appearance like that of water. Thus, it would apply to XYZ as much as to H₂O.

It follows from my reflections on Laurence and Margolis that this way of dealing with the problem of accounting for the content of our concepts in the light of Putnam's Twin Earth thought experiment is problematic with the upshot that Prinz cannot explain how Oscar and Twin Oscar can diverge in their concepts. Given that the twins also acquire concepts that apply to kinds that do not have microphysical essences (for example, MILK) they cannot acquire concepts such as WATER or TWIN WATER by using an unarticulated notion of essence. Rather, they must each have a range of different notions of essence and apply the appropriate one in each particular case. For example, Oscar must apply a notion of microphysical essence when interacting with water and a notion of biofunctional essence when interacting with milk otherwise he will not acquire the concepts WATER and MILK on the basis of those respective interactions. But, once again, this raises the question of how Oscar could have acquired such notions of essence and know which to apply in each particular case in the absence of explicit instruction. In sum then, Prinz's invocation of psychological essentialism does not help explain how the Putnamian Twins, and by extension us, acquire the concepts that they acquire. Thus, the problem that I directed at Laurence and Margolis is not localised to a Fodorian LOT view of concepts.

6. Conclusion

Psychological essentialism implies that many of our concepts are not appearance concepts; rather, we think of them as picking out categories whose instances are

bound together by their hidden properties. Both Laurence and Margolis and Prinz have appealed to such essentialist commitments to explain how it is possible for us to learn natural kind concepts. I have argued that matters are not so straightforward. Although Fodor is mistaken in arguing that it is impossible to learn primitive concepts, a commitment to essentialism will only help us to reliably learn concepts if we have a range of sufficiently articulated notions of essences and that we apply the appropriate notion in each particular case. Given that essences come in different forms, a one-size-fits-all notion of essence will not enable us to acquire a full range of concepts including WATER, MILK, DOG and RAT. I have also expressed scepticism that a typical child's experiences are rich enough to provide her with such a range of notions of essence or with reliable evidence as to which applies in each particular case.

In the light of this what conclusion should be drawn about the role of learning in concept acquisition? There are three possibilities. The first is to concede that although we do not currently have an adequate account of how we reliably learn concepts such as WATER, MILK, DOG, and so on, the idea that we reliably learn such concepts is non-negotiable so that we should continue to search for an adequate learning theory. The second is to take a nativist view of such concepts seriously: perhaps concepts such as WATER, MILK, and DOG are part of our innate endowment. Taking this line involves regarding my reflections as providing a poverty of the stimulus argument for nativism, an argument to the effect that a typical individual's experiences are not rich enough to facilitate the reliable learning of many of the concepts that she acquires so that those concepts must be innate if they are to be shared. To adopt such a point of view is not to endorse a form of nativism as radical as that championed by Fodor (1975), (1981). For, one could consistently argue that to view, say, WATER and

MILK as innate in no way commits one to the view that concepts such as CARBURETOR are innate. For, it is plausible that concepts such as the latter are acquired on the basis of a rich battery of explicit instruction that is typically absent in the case of the former. The third option would be to reject both the views that concepts such as WATER, MILK and DOG are innate and that they are reliably learned. On this view we do acquire concepts on the basis of our experiences and the concepts that we so acquire are not innate. However, there is an accidental or random element involved in the concept acquisition process that is overlooked by the orthodox pro-learning view. For example, children attempting to learn the concept expressed by the word 'water' may well assume that the adult members of their community all express the same concept by means of that word and aim to pick up that very concept. However, given that they are presented with impoverished evidence as to the nature of that concept, would-be learners are in effect forced to make leaps into the dark. The upshot of this is that different children make leaps in different directions and so acquire different concepts. Thus, many children will not have acquired the 'correct' concept and those that have will have done so by accident. In the absence of subsequent experiences (for, example, in the form of explicit instruction as to the nature of the category picked out by the target concept) to correct mistakes or substantiate lucky guesses, such variance in the concepts expressed by a given word across the community will have the result that there is no univocal concept expressed by that word. The upshot of this is that the very assumptions that underpin the orthodox view that we learn our concepts will not hold. These are the assumptions that concepts are widely shared across communities and that there are reliable mechanisms for ensuring that individuals (particularly children) being

inducted into a given community acquire the concepts employed by its mature members.

I will not argue for any one of these conclusions in preference to its competitors but close by noting that all of them are noteworthy and present a challenge to the orthodox view of the role of learning in concept acquisition.

References

Block, N. (1986). Advertisement for a semantics for psychology. In P.A. French, T.E. Ueling Jr. and H.K. Wettstein (Eds.), *Midwest Studies in Philosophy, X: Studies in the Philosophy of Mind*. Minneapolis, MN: University of Minnesota Press.

Bloom, P. (2000). *How Children Learn the Meaning of Words*. Cambridge, MA: MIT Press.

Bloom, P. (2004). *Descartes' Baby: How the Science of Child Development Explains What Makes us Human*. New York: Basic Books.

Bloom, P. (2010). *How Pleasure Works: The New Science of Why we Like what we Like*. London: Bodley Head.

Boeckz, C. (2010). *Language in Cognition*. Oxford: Wiley-Blackwell.

Carey, S. (2009). *The Origin of Concepts*. New York: Oxford University Press.

Chomsky, N. (1986). *Knowledge of Language*. Westport, CT: Praeger.

Chomsky, N. (2009). *Cartesian Linguistics: A Chapter in the History of Rationalist Thought, 3rd Edition*. Cambridge: Cambridge University Press.

Churchland, P.S. (1986). *Neurophilosophy*. Cambridge, MA: MIT Press.

Cowie, F. (1999). *What's Within: Nativism Reconsidered*. New York: Oxford University Press.

Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.

Fodor, J. (1975). *The Language of Thought*. Cambridge, Mass: Harvard University Press.

Fodor, J. (1981). The present status of the innateness controversy. In his *Representations*. Cambridge, Mass: MIT Press.

Fodor, J. (1987). *Psychosemantics*. Cambridge, Mass: MIT Press.

Fodor, J. (1990) . *A Theory of Content and Other Essays*. Cambridge, Mass: MIT Press.

Fodor, J. (1998). *Concepts*. Oxford: Oxford University Press.

Fodor, J. (2001). Doing without what's within. *Mind*, 110, 99-148.

Fodor, J. (2008) *LOT 2: The Language of Thought Revisited*. Oxford: Oxford University Press.

Fodor, J. and LePore, E. (2002). *The Compositionality Papers*. Oxford: Oxford University Press.

Gelman, S. (2003). *The Essential Child*. New York: Oxford University Press.

Gelman, S. (2004). Psychological essentialism in children. *Trends in Cognitive Sciences*, 8, 404-409.

Isac, D. and Reiss, C. (2013). *I-Language: An Introduction to Linguistics as Cognitive Science*, 2nd Edition. Oxford: Oxford University Press.

Jackendoff, R. (1989). What is a concept that a person may grasp it? *Mind and Language*, 4, 68-102.

Keil, F. (1989) *Concepts, Kinds, and Cognitive Development*. Cambridge, MA: MIT Press.

Kripke, S. (1980). *Naming and Necessity*. Oxford: Blackwell.

Laurence, S. and Margolis, E. (2002). Radical concept nativism. *Cognition*, 86, 25-55.

Machery, E. (2009). *Doing Without Concepts*. Oxford: Oxford University Press.

Macnamara, J. (1986). *Border Dispute: The Place of Logic in Psychology*.

Cambridge, MA: MIT Press.

Mandler, J. (2004). *The Foundations of Mind*. New York: Oxford University Press.

Mandler, J. (2008). On the birth and growth of concepts. *Philosophical Psychology*, 21, 207-230.

Margolis, E. (1998). How to acquire a concept. *Mind and Language*, 13, 347-369.

Margolis, E. and Laurence, S. (2011). Learning matters: the role of learning in concept acquisition. *Mind and Language*, 22, 507-539.

Medin, D. and Ortony, A. (1989). Psychological essentialism. In S. Vosniadou (Ed.) *Similarity and Analogical Reasoning*. New York: Cambridge University Press.

Murphy, D. (2002). *The Big Book of Concepts*. Cambridge, Mass: MIT Press.

Pinker, S. (1997). *How The Mind Works*. New York: W.W. Norton.

Pinker, S. (2007). *The Stuff of Thought*. London: Alan Lane.

Prinz, J. (2002). *Furnishing the Mind: Concepts and their Perceptual Basis*.

Cambridge, MA: MIT Press.

Prinz, J. (2005) The Return of concept empiricism. In H. Cohen and C. Lefebvre

(Eds.) *Categorization and Cognitive Science*. Amsterdam: Elsevier Science.

Putnam, H. (1975). The meaning of 'meaning'. In his *Mind, Language and Reality:*

Philosophical Papers Volume 2. Cambridge: Cambridge University Press.

Samet, J. and Flanagan, O. (1989). Innate representations. In S. Silvers (Ed.)

Representations. New York: Kluwer.

Soja, N., Carey, S. and Spelke, E. (1991). Ontological categories guide young

children's inductions on word meaning: object terms and substance terms. *Cognition*,

38, 179-211.

Sterelny, K. (2003). *Thought in a Hostile World: The Evolution of Human Cognition*.

Oxford: Blackwell Publishing.

Weiskopf, D. (2008). The origins of concepts. *Philosophical Studies*, 140, 359-384.